arXiv:1703.10959v2 [cs.DC] 3 Apr 2017

DRAFT - Parallelism, Concurreny and Distribution in Constraint Handling Rules: A Survey

THOM FRÜHWIRTH
University of Ulm, Germany

April 4, 2017

Abstract

Constraint Handling Rules (CHR) is both an effective concurrent declarative programming language and a versatile computational logic formalism. CHR has its roots in constraint logic programming and concurrent constraint programming, but also integrates ideas from multiset transformation and rewriting systems. CHR programs consist of guarded reactive rules that transform multisets of constraints. One of the main features of CHR is its inherent concurrency. Intuitively, rules can be applied to parts of a multiset in parallel.

In this comprehensive survey, we give an overview of concurrent and parallel as well as distributed CHR semantics, standard and more exotic, that have been proposed over the years at various levels of refinement. These semantics range from the abstract to the concrete. They are related by formal soundness results. Their correctness is established as a correspondence between parallel and sequential computations.

On the more practical side, we present common concise sample CHR programs that have been widely used in experiments and benchmarks. We review parallel CHR implementations in software as well as hardware. The experimental results obtained show a consistent parallel speed-up. Most implementations are available online.

Due to its high level of abstraction, the CHR formalism can also be used to implement and reason with models for concurrency. To this end, the Software Transaction Model, the Actor Model, Colored Petri Nets and the Join-Calculus have been faithfully encoded in CHR. Finally, we summarize the state of art and indicate what problems are left open for future research.

KEYWORDS: Parallelism, Concurrency, Distribution, Constraint Handling Rules,
Constraint Logic Programming, Concurrent Constraint Programming,
Semantics, Rewriting, Concurrency Models.

Contents

1 Introduction **2**

2 Parallel Abstract Operational Semantics of CHR **5**
2.1 Semantics of CHR and their Properties 5
2.2 Abstract Syntax of CHR 5
2.3 Sequential Abstract Operational Semantics of CHR 6
2.4 Extension to Parallel Abstract Semantics 7
2.5 Monotonicity, Soundness and Serializability 7

3 Parallel CHR Example Programs **8**
3.1 Algorithms of Erastothenes, Euclid, von Neumann, Floyd-Warshall 8
3.2 Classical Models and Classical Algorithms with Statefulness 9
3.3 Parallel Preflow-Push 9
3.4 Parallelizing the Union-Find Algorithm 10

4 Parallel CHR with Transactions **11**
4.1 Transactions in Parallel CHR . . . 11
4.2 Abstract Syntax and Semantics of CHRt . . . 11
4.3 Monotonicity, Soundness and Serializability . . . 12
4.4 Encoding Transactions in Standard CHR . . . 12

5 A More Refined Parallel Semantics **12**
5.1 Syntax for Refined Parallel CHR . . . 13
5.2 Sequential Refined CHR Semantics . . . 13
5.3 Extension to Parallel Refined CHR Semantics . . . 14
5.4 Monotonicity, Soundness and Serializability . . . 14

6 Parallel CHR Implementation in Haskell **15**
6.1 Implementation Principles . . . 15
6.2 Experimental Results . . . 16

7 Massively-Parallel Set-Based CHR Semantics **17**
7.1 Massively-Parallel Set-Based Semantics . . . 17
7.2 Example Programs under Exhaustive Parallelism . . . 18
7.3 Soundness and Deletion-Acyclicity . . . 19

8 Parallel Hardware Implementations **20**
8.1 Basic Compilation of CHR to Procedural Languages . . . 20
8.2 Compiling CHR to Parallel FPGA Hardware via VHDL . . . 20
8.3 Parallel Execution on a Graphical Processing Unit GPU . . . 22

9 Distribution in CHR **22**
9.1 Distributed Set-Based Goal Stores in CHRd . . . 23
9.2 Distributed Parallel CHRe . . . 24
9.3 Syntax of CHRe . . . 24
9.4 Abstract and Refined Semantics of CHRe . . . 24
9.5 Encoding 1- and n-Neighbor Rules in Local Rules . . . 26

10 Models of Concurrency in CHR **27**
10.1 Software Transactional Memory STM . . . 27
10.2 Colored Petri Nets CPN . . . 29
10.3 Actor Model . . . 29
10.4 Join-Calculus and Join-Patterns . . . 30

11 Conclusions **31**

1 Introduction

Before we start with our survey, we shortly clarify the essential notions at stake and introduce Constraint Handling Rules. The technical terms of concurrency, parallelism and distribution have an overlapping meaning, and processes are another central notion in this context:

Concurrency allows for logically more or less independent computations, be the sequential or parallel. This abstract notion thus supports the modular design of independent components that can be composed together.

Parallelism allows for computations that happen simultaneously, thus hopefully improving performance. With the arrival of multi-core processors, it has become a dominant computation model. The processors may have access to a shared memory to exchange information.

Distribution allows for program components that are located on physically distributed decentralized networked processors. Each processor has its own local memory (distributed memory).

Processes are programs that are executed independently but can interact with each other. Processes can either execute local actions or communicate, coordinate and synchronize by passing (sending and receiving) messages. Depending on context and level of abstraction, processes are also called threads, workers, tasks, activities or even agents.

Concurrency and distribution are easier with declarative programming languages, since they are compositional: different computations can be composed into one without unintended interference. Moreover, declarative languages offer a wealth of program analysis and reasoning techniques.

Constraint Handling Rules (CHR). CHR is both an effective concurrent declarative constraint-based programming language and a versatile computational logic formalism [Frü09, SVWSDK10, FR11, Frü15, Frü16]. While conceptually simple, CHR is distinguished by a remarkable combination of desirable features:

- a semantic foundation in classical logic as well as in linear logic,
- an effective and efficient sequential and parallel execution model,
- a proof that every algorithm can be expressed with best known time and space complexity,
- up to a million rule applications per second for CHR implemented in C or hardware,
- guaranteed properties like the anytime algorithm and online algorithm properties,
- program analysis methods for deciding essential properties like confluence and program equivalence.

Information on CHR can be found online at `http://www.constraint-handling-rules.org`, including news, tutorials, papers, bibliography, online demos and free downloads of the language.

Minimum Example. Assume we would like to compute the minimum of some numbers, given as multiset `min(`n_1`)`, `min(`n_2`)`,..., `min(`n_k`)`. We interpret the constraint (predicate) `min(`n_i`)` to mean that the number n_i is a candidate for the minimum value. We make use of the following CHR rule that filters the candidates.

```
min(N) \ min(M) <=> N=<M | true.
```

The rule consists of a left-hand side, on which a pair of constraints has to be matched, a guard check `N=<M` that has to be satisfied, and an empty right-hand side denoted by `true`. In effect, the rule takes two `min` candidates and removes the one with the larger value (constraints after the \ symbol are to be removed). Starting with a given initial state, CHR rules are applied exhaustively, resulting in a final state. Note that CHR is a committed-choice language, i.e. there is no backtracking in the rule applications. Here the rule keeps on going until only one, thus the smallest value, remains as single `min` constraint.

A *state* is a multiset of constraints. In a sequential computation, we apply one rule at a time to a given state. A possible computation sequence is (where we underline constraints involved in a rule application):

$$\begin{array}{c} \underline{\texttt{min(1)}}\texttt{, }\underline{\texttt{min(0)}}\texttt{, min(2), min(1)} \mapsto \\ \underline{\texttt{min(0)}}\texttt{, }\underline{\texttt{min(2)}}\texttt{, min(1)} \mapsto \\ \underline{\texttt{min(0)}}\texttt{, }\underline{\texttt{min(1)}} \mapsto \\ \texttt{min(0)} \end{array}$$

The final state is called *answer*. The remaining constraint contains the minimum value, in this case zero.

By the way, CHR insists on multisets so one can directly model resources as constraints, for example:

```
buy : cup \ euro, euro <=> coffee.
```

This rule expresses that we get a coffee for two euros if we have a cup.

Concurrency and Parallelism in CHR. One of the main features of CHR is its inherent concurrency. Intuitively, in a parallel execution of a CHR program, rules can be applied to separate parts of a state in parallel. As we will see, CHR rules even even be applied in parallel to overlapping parts of a state, without the need to change the program. This is referred to as *logical parallelism* or *declarative concurrency*.

The rule of `min` can be applied in parallel to different parts of the state:

$$\begin{array}{ccc} \underline{\texttt{min(1)}},\ \underline{\texttt{min(0)}}, & & \underline{\texttt{min(2)}},\ \underline{\texttt{min(1)}} \mapsto \\ \underline{\texttt{min(0)}}, & & \underline{\texttt{min(1)}} \mapsto \\ & \texttt{min(0)} & \end{array}$$

We arrive at the answer in less computation steps than with the sequential execution.

The rule can also be applied in parallel to overlapping parts of the state, provided the overlap is not removed by any rule. For example, let the overlap be the constraint `min(0)`. Then the three pairs `min(0), min(1)`,`min(0), min(1)` and `min(0), min(2)` can be matched to different rule instances. (Note that we always match the same `min(0)`, but that we have two copies of `min(1)`.) These rules can be applied at the same time, since the common (overlapping) constraint `min(0)` is not removed.

$$\begin{array}{cc} \underline{\underline{\texttt{min(0)}}}, & \underline{\texttt{min(1)}},\ \underline{\texttt{min(2)}},\ \underline{\texttt{min(1)}} \mapsto \\ & \texttt{min(0)} \end{array}$$

So this is another, even shorter way to arrive at the same answer.

In CHR, concurrently executing processes are CHR constraints that communicate via a shared built-in constraint store. The built-in constraints take the role of (partial) messages and variables take the role of communication channels.

Guaranteed Properties of CHR. We can stop the computation at *any time* and observe the current state as intermediate answer. We can then continue by applying rules to this state without the need to recompute from scratch. If we stop again, we will observe another intermediate answer that is closer to the final answer. The intermediate answers more and more approximate the final answer. This property of a CHR program is called the *anytime algorithm property*. Note that by this description, an anytime algorithm is also an *approximation algorithm*.

Now assume that while the program runs, we add another constraint. It will eventually participate in the computation in that a rule will be applied to it. The answer will be as if the newly added constraint had been there from the beginning but ignored for some time. This property of a CHR program is called *incrementality* or *online algorithm property*.

Last but not least, the essential *monotonicity property* of CHR means that adding constraints to a state cannot inhibit the applicability of a rule. (Rule matching and guards check for presence of certain constraints, never absence.) Among other things, this monotonicity enables decidable program analyses and helps declarative concurreny. Most, but not all semantics that we introduce enjoy the monotonicty property.

Confluence. This property of a program guarantees that any computation starting from a given initial state results in the same answer no matter which of the applicable rules are applied. There is a decidable, sufficient and necessary syntactic condition to analyse confluence of terminating programs and to detect rule pairs that lead to non-confluence when applied. Among other things, confluence implies that rules can be applied in parallel with the same result as any sequential computation.

Overview of the paper. In the next section we define abstract syntax and abstract operational semantics for CHR. One sequential transition describes rule applications, another one parallel parallel transitions, a third one connects the two. In section 3, we present common classic CHR example programs based on well-known algorithms. Often one rule suffices. All but one of the programs can be run in parallel without change.

In section 4, we extend parallel CHR with transactions, a popular and essential concept in concurrency. In section 5, we refine our abstract semantics by differentiating between a goal store and a constraint store. The goal store holds active constraints that execute as processes, the constraint store holds suspended constraints as data. In section 6, we describe an implementation of the refined semantics in Haskell.

Section 7 introduces a more exotic set-based abstract semantics that is massively parallel. In the extreme case, this theoretical model allows to sort in constant time and to solve SAT problems in linear time. We then move on to more mundane fast hardware implementations of CHR in section 8. In section 9 we discuss distributed semantics for CHR,

where the constraint store and computations are decentralized. Last but not least, in section 10 we show the benefit of embedding common models of concurrency in CHR on four concrete models. We end the paper with conclusions and directions for future work.

For a better reading experience, we use the editorial we throughout. Of course it refers to different authors in different sections of this paper.

2 Parallel Abstract Operational Semantics of CHR

We will present the sequential equivalence-based abstract CHR semantics and extend it with parallelism. We assume basic familiarity with first-order predicate logic and state transition systems. Readers familiar with CHR can skip this section. We start with some preliminaries.

2.1 Semantics of CHR and their Properties

Structural Operational Semantics (SOS) is a common inductive approach to describe the behavior of programming languages, in particular concurrent ones. In SOS, a *state transition system* specifies the computations. Transitions rewrite states and take the form of inference rules. All semantics of CHR, sequential or parallel, employ this approach.

Semantics for sequential CHR. They exist in various formulations and at various levels of refinement, going from the abstract to the concrete (refined) [Frü09, BRF10a]:

- The *very abstract semantics* [Frü09] is close to modus ponens of predicate logic.
- The *abstract semantics* [AFM99] is the classical basis for CHR program analysis and its properties.
- The more recent *state-equivalence-based abstract semantics* [RBF09] will be the starting point of our survey. We will extend it with parallelism.
- The *refined semantics* [DSGH04] describes more concretely the actual behavior of CHR implementations. All more concrete parallel semantics of CHR are based on it.

In addition, several alternative operational semantics for sequential CHR have been proposed, among them [GM09, Hae11, BRF10b].

Soundness and Serializability. The *correctness* of a more refined semantics is shown by its *soundness* with regard to a more abstract semantics. This means that for each computation in the refined semantics, there is a corresponding computation in the abstract semantics. The converse (completeness) typically does not hold, because refined semantics are more concrete and thus rule out certain computations. When we introduce a parallel semantics for CHR, it will be related by soundness to a more abstract semantics and/or the sequential part of the semantics.

Actually, the *interleaving semantics* approach to concurrency is defined by the fact that for each possible parallel computation, there exists a corresponding sequential computation with the same result. The sequential computation uses interleaving of the different parallel computations. This means that a parallel computation step can be simulated by a sequence of sequential computation steps. This correspondence property is called *serializability (sequential consistency)*. Most semantics we discuss are correct in this way.

2.2 Abstract Syntax of CHR

Constraints are relations, distinguished predicates of first-order predicate logic. We differentiate between two kinds of constraints: *built-in (pre-defined) constraints* and *user-defined (CHR) constraints* which are defined by the rules in a CHR program. Built-in constraints can be used as tests in the guard as well as for auxiliary computations in the body of a rule. In this survey, besides the trivial constraint `true`, we will have syntactical equality = between logical terms and equations between arithmetic expressions.

Definition 2.1. A *CHR program* is a finite set of rules. A *(generalized) simpagation rule* is of the form

$$r : H_1 \backslash H_2 \Leftrightarrow C|B$$

where r : is an optional *name* (a unique identifier) of a rule. In the rule *head* (left-hand side), H_1 and H_2 are conjunctions of user-defined constraints, the optional *guard* $C|$ is a conjunction of built-in constraints, and the *body* (right-hand side) B is a goal. A *goal* is a conjunction of built-in and user-defined constraints. A *state* is also a goal. Conjunctions are understood as *multisets* of their conjuncts.

In the rule, H_1 are called the *kept constraints*, while H_2 are called the *removed constraints*. At least one of H_1 and H_2 must be non-empty. If H_1 is empty, the rule corresponds to a simplification rule, also written

$$s : H_2 \Leftrightarrow C|B.$$

If H_2 is empty, the rule corresponds to a propagation rule, also written

$$p : H_1 \Rightarrow C|B.$$

Interestingly, most parallel semantics we consider do not allow for propagation rules.

Ground CHR. Many semantics assume that variables are substituted by ground (variable-free) terms at run-time. This requirement can be captured by a common syntactic fragment of CHR. *Ground CHR* is a syntactic fragment of CHR where every variable in a rule (also) occurs in the head of the rule. We also say that the rule is *range-restricted*. So given a ground initial states, all states in a computation will stay ground. It is worth noting that Ground CHR without propagation rules is still Turing-complete: it can implement a Turing machine with just one rule as we will see in section 3.2.

2.3 Sequential Abstract Operational Semantics of CHR

The semantics follows [RBF09, Bet14]. It relies on a structural equivalence between states that abstracts away from technical details in a transition.

State Equivalence. The equivalence relation treats built-in constraints semantically and user-defined constraints syntactically. Basically, two states are equivalent if they are logically equivalent (imply each other) while taking into account that user-defined constraints form a multiset, i.e. multiplicities matter. For a state S, the notation S_{bi} denotes the built-in constraints of S and S_{ud} denotes the user-defined constraints of S.

Definition 2.2 (State Equivalence)**.** Two states $S_1 = (S_{1bi} \wedge S_{1ud})$ and $S_2 = (S_{2bi} \wedge S_{2ud})$ are *equivalent*, written $S_1 \equiv S_2$, if and only if

$$\models \forall(S_{1bi} \rightarrow \exists\bar{y}((S_{1ud} = S_{2ud}) \wedge S_{2bi})) \wedge \forall(S_{2bi} \rightarrow \exists\bar{x}((S_{1ud} = S_{2ud}) \wedge S_{1bi}))$$

with $\bar{x}$ those variables that only occur in S_1 and $\bar{y}$ those variables that only occur in S_2.

The CHR state equivalence is defined by two symmetric implications and moreover syntactically equates the conjunctions of user-defined constraints as multisets. For example,

$$X{=}{<}Y \wedge Y{=}{<}X \wedge c(X,Y) \;\equiv\; X{=}Y \wedge c(X,X) \not\equiv X{=}Y \wedge c(X,X) \wedge c(X,X).$$

Transition. Using this state equivalence, the abstract CHR semantics is defined by a single transition that is the workhorse of CHR program execution. It defines the application of a rule. Let the rule $(r : H_1 \backslash H_2 \Leftrightarrow C|B)$ be a variant of a rule from a given program $\mathscr{P}$. A variant (renaming) of an expression is obtained by uniformly replacing its variables by fresh variables.

$$\text{(Apply)} \quad \frac{S \equiv (H_1 \wedge H_2 \wedge C \wedge G) \quad (r : H_1 \backslash H_2 \Leftrightarrow C|B) \in \mathscr{P} \qquad (H_1 \wedge C \wedge B \wedge G) \equiv T}{S \mapsto_r T}$$

Upper-case letters stand for (possibly empty) conjunctions of constraints in this section. The goal G is called *context* of the rule application. It is left unchanged.

In a *transition (computation step)* $S \mapsto_r T$, S is called *source state* and T is called *target state*. We may drop the reference to the program $\mathscr{P}$ and rule r to simplify the presentation.

If the source state can be made equivalent to a state that contains the head constraints and the guard built-in constraints of a variant of a rule, then we delete the removed head constraints from the state and add the rule body constraints to it. Any state that is equivalent to this target state is in the transition relation.

A *computation (derivation)* of a goal S in a program P is a connected sequence $S_i \mapsto S_{i+1}$ beginning with the *initial state (query)* S_0 that is S and ending in a *final state (answer, result)* or the sequence is *non-terminating (diverging)*. The notation $\mapsto^*$ denotes the reflexive and transitive closure of $\mapsto$.

For the minimum example, here is a possible (Apply) transition from a state $S = (min(0) \wedge min(2) \wedge min(1))$ to a state $T = (min(0) \wedge min(1))$:

$$\frac{\begin{array}{c} S \equiv (min(X) \wedge min(Y) \wedge X \leq Y \wedge (X = 0 \wedge Y = 2 \wedge min(1))) \\ (min(X) \backslash min(Y) \Leftrightarrow X \leq Y | true) \\ (min(X) \wedge X \leq Y \wedge true \wedge (X = 0 \wedge Y = 2 \wedge min(1))) \equiv T \end{array}}{S \mapsto T}$$

Note that the abstract semantics does not account for termination of propagation rules: If a state can fire a propagation rule once, it can do so again and again, ad infinitum. This is called trivial non-termination of propagation rules. Most parallel semantics rule out propagation rules. They will be discussed for distributed CHR, though.

2.4 Extension to Parallel Abstract Semantics

We extend the abstract semantics by parallelism. We interpret conjunction as parallel operator. As we have seen for the minimum example, CHR rules can also be applied simultaneously to overlapping parts of a state, as long as the *overlap* (shared, common part) is not removed by any rule. Following [Frü05a], CHR parallelism with overlaps is called *strong*. It can be defined as follows, see also chapter 4 in [Frü09].

Definition 2.3 ((Strong) Parallelism (with Overlap))**.** We denote parallel transitions by the relation $\Mapsto$.

$$\text{(Intro-Par)} \quad \frac{A \mapsto C}{A \Mapsto C}$$

$$\text{(Parallel)} \quad \frac{A \wedge E \Mapsto C \wedge E \qquad B \wedge E \Mapsto D \wedge E}{A \wedge B \wedge E \Mapsto C \wedge D \wedge E}$$

The transition (Intro-Par) says that any sequential transition is also a parallel transition. The transition (Parallel) combines two parallel transitions using conjunction into a single parallel transition where the overlap E is kept.

Again, back to the minimum example:

$$\text{(Parallel)} \quad \frac{min(1) \wedge min(0) \Mapsto true \wedge min(0) \qquad min(2) \wedge min(0) \Mapsto true \wedge min(0)}{min(1) \wedge min(2) \wedge min(0) \Mapsto true \wedge true \wedge min(0)}$$

2.5 Monotonicity, Soundness and Serializability

The *monotonicity property* of CHR states that adding constraints to a state cannot inhibit the applicability of a rule. It is easy to see from the context of the sequential (Apply) transition and from the overlap of the (Parallel) transition that a rule can be applied in any state that contains its head and guard.

Theorem 2.1 (Monotonicity of CHR)**.** If $A \mapsto B$ then $A \wedge G \mapsto B \wedge G$. If $A \Mapsto B$ then $A \wedge E \Mapsto B \wedge E$.

The *correctness* of the abstract parallel semantics can be established by proving the following theorem.

Theorem 2.2 (Soundness and Serializability)**.** If $A \Rightarrow B$ in parallel, then there exists a sequential computation $A \mapsto^* B$.

The essential aspect of the truth is that the (Parallel) transition can be simulated sequentially: If $A \wedge E \mapsto B \wedge E$ and $C \wedge E \mapsto D \wedge E$, then $A \wedge C \wedge E \mapsto S \mapsto B \wedge D \wedge E$, where S is either $A \wedge D \wedge E$ or $B \wedge C \wedge E$, i.e. the two transitions commute.

3 Parallel CHR Example Programs

These exemplary CHR programs are mostly folklore in the CHR community, see e.g. chapters 2 and 7 in [Frü09]. These are concise and effective implementations of classical algorithms starting with prime numbers and sorting and ending with Preflow-Push and Union-Find. Often one type of constraint and one rule will suffice, and we will not need more then six rules. Due to the guaranteed properties of CHR, these implementations are also incremental anytime online approximation algorithms and they run in parallel without any need for modifying the program. An exception is Union-Find, which is known to be hard to parallelize. We do it with the help of confluence analysis. These example programs have been used in experiments in most papers surveyed here. Experimental results consistently report parallel speed-ups for these problems.

3.1 Algorithms of Erastothenes, Euclid, von Neumann, Floyd-Warshall

Here we introduce some classical algorithms over numbers and graphs. They are implemented as simple multiset transformations reminiscent of the Chemical Abstract Machine (CHAM). Typically, they can be implemented with one kind of constraint and a single rule in CHR that can be applied in parallel to pairs of constraints. Our running example of minimum falls into this category. Correctness of each implementation can be shown by contradiction: given the specified initial goal, if the resulting answer were not of the desired form, the rule would still be applicable.

Prime Numbers. The following rule is like the rule for minimum, but the guard is different, more strict. In effect, it filters out multiples of numbers, similar to the Sieve of Erastothenes.

```
sift : prime(I) \ prime(J) <=> J mod I =:= 0 | true.
```

If all natural numbers from 2 to n are given, only the prime numbers within this range remain, since non-prime numbers are multiples of other numbers greater equal to 2.

Greatest Common Divisor (GCD). The following rule computes the greatest common divisor of natural numbers written each as `gcd(N)`.

```
gcd(N) \ gcd(M) <=> 0<N,N=<M | gcd(M-N).
```

The rule replaces M by the smaller number $M-N$ as in Euclid's algorithm. The rule maintains the invariant that the numbers have the same greatest common divisor. Eventually, if $N = M$, a zero is produced. The remaining nonzero `gcd` constraint contains the value of the gcd.

Merge Sort. The initial goal state contains arcs of the form `a->V` for each value `V`, where `a` is a given smallest (dummy) value.

```
msort : A->B \ A->C <=> A<B, B<C | B->C.
```

The rule only updates the first argument of the `arc` constraint, never the second. The first argument is replaced by a larger value and the two resulting arcs form a small chain `A->B, B->C`. The rule maintains the invariant that `A=<B`. So eventually, in each arc, a number will be followed by its immediate successor, and thus the resulting chain of arcs is sorted.

For optimal run-time complexity, we prefer merging arc chains of the same length. To this end, we precede each chain with its length, written `N=>FirstNode`. We also have to add a rule to initiate merging of chains of the same length:

```
N=>A, N=>B <=> A<B | N+N=>A, A->B.
```

In the initial goal we now introduce constraints of the form `1=>V` for each value `V`.

Floyd-Warshall All-Pair Shortest Paths. Our implementation finds the shortest distance between all connected pairs of nodes in the transitive closure of a directed graph whose edges are annotated with non-negative distances.

```
shorten : arc(I,K,D1), arc(K,J,D2) \ arc(I,J,D3) <=>
                              D3>D1+D2 | arc(I,J,D1+D2).
```

Clearly we can shorten arc distances in parallel.

3.2 Classical Models and Classical Algorithms with Statefulness

These algorithms about abstract problems are characterized by their *statefulness*, i.e. their essence is a state change, an update.

Turing Machine. The Turing machine is the classical model of computability.

```
st(QI,SI,SJ,D,QJ) \ state(I,QI), cell(I,SI) <=> state(I+D,QJ), cell(I,SJ).
```

The state transition steps of the Turing machine are given as constraints `st(QI,SI,SJ,D,QJ)`: in the current state `QI` reading tape symbol `SI`, write symbol `SJ` and move in direction `D` to be in state `QJ`. The direction is either left or right, we move along the cells of a tape. We represent cells as an array, so positions are numbers and the direction is either $+1$ or -1. A Turing machine with one tape is inherently sequential, since we can only be in one state at a time.

The implementation of the Turing machine shows Turing-completeness of the Ground CHR fragment with constants only and without propagation rules, actually with a single rule [Sne08].

Dijkstras Dining Philosophers. In this classical problem in concurrency, several philosophers sit at a round table. Between each of them a fork is placed. A philosopher either thinks or eats. In order to eat, a philosopher needs two forks, the one from his left and the one from his right. After a while, an eating philosopher will start to think again, releasing the forks and thus making them available to his neighbors again.

```
think_eat : think(X), fork(X), fork(Y) <=> Y =:= (X+1) mod n | eat(X).
eat_think : eat(X) <=> Y =:= (X+1) mod n | think(X), fork(X), fork(Y).
```

In the implementation, we assume a given number `n` of philosophers (and forks). They are identified by a number from zero to `n-1`. The rules are inverses of each other, the constraints simply switch sides.

The problem is to design a concurrent algorithm that is fair, i.e. that no philosopher will starve. Here we are mainly interested in the inherent parallelism of the problem. For the experiments, time counters for eating and thinking were introduced.

Blockworld. Blockworld simulates robot arms re-arranging stacks of blocks.

```
grab  : grab(R,X),  empty(R), clear(X), on(X,Y) <=> hold(R,X), clear(Y).
putOn : putOn(R,Y), hold(R,X), clear(Y) <=> empty(R), clear(X), on(X,Y).
```

The *operation constraints* `grab` and `putOn` specify the action that is taken. The other constraints are *data constraints* holding some information about the scenario. The rule `grab` specifies that robot arm `R` grabs block `X` if `R` is empty and block `X` is clear on top and on block `Y`. As a results, robot arm `R` holds block `X` and block `Y` is clear. The rule `putOn` specifies the inverse action. The data constraints in the rule switch sides. Parallelism is induced by several robot arms and several actions present.

3.3 Parallel Preflow-Push

A flow network is a directed graph, where each edge is assigned a non-negative capacity. We want to find a maximum flow through the network from a source to a sink node under the capacity restrictions. The Preflow-Push algorithm moves flow locally between neighboring nodes until a maximum flow is reached.

Parallel implementations of the preflow-push algorithm are usually low-level. In [Mei07], we present and analyse a concise parallel implementation of the preflow-push algorithm by just four rules. Due to lack of space we cannot explain these algorithms in any detail. In the code listing below, comment lines start with the symbol `%`.

```
 % increase node height by one, remove minimum
lift : n(U,N), e(U,E) \ h(U,_), m(U,M,C)
       <=> U \= source, U \= sink, 0 < E, C =:= N+E | h(U,M+1).
 % replace K by HU in unchecked egde, insert minimum
up   : h(U,HU), h(V,HV) \ r(U,V,K)
       <=> HU =< HV, K < HU | m(U,HV,1), r(U,V,HU).
 % push flow downwards by one unit, insert minimum, reverse edge
push : h(U,HU), h(V,HV) \ e(U,EU), e(V,EV), r(U,V,_)
       <=>  0 < EU, HV < HU | e(U,EU-1), e(V,EV+1), m(V,HU,1), r(V,U,HV).
 % compute minimum for node, count for completeness
min  : m(U,M1,C1), m(U,M2,C2) <=> m(U,min(M1,M2),C1+C2).
```

The variable `U` stands for a node, `N` is its number of outward capacity edges, `E` is its current excess flow, `HU` is its current height. The constraint `m(U,M,C)` encodes a minimum candidate `M` for `U`, where the counter `C` allows to detect if the minimum of all outward edges has been computed. The constraint `r(U,V,K)` encodes a residual edge from `U` to `V` with remaining capacity `K`.

The implementation described in [Mei07] simulates parallel computations sequentially using an interleaving semantics approach and time stamps for user-defined constraints. We greedily, randomly and exhaustively apply rules for a given time t before progressing to time $t+1$. The consistent speed-up observed in experiments with random graphs depends on the total amount of flow units, its distribution on disjoint nodes, and the density of the flow network.

3.4 Parallelizing the Union-Find Algorithm

The classical union-find (also: disjoint set union) (UF) algorithm efficiently maintains disjoint sets under the operation of union. Each set is represented by a rooted tree, whose nodes are the elements of the set. Union-Find is acknowledged in the literature to be hard to parallelize.

In [Frü05a], we implement the UF algorithm with optimal time and space complexity in CHR (with the anytime online algorithm properties). This effectiveness is believed impossible in other pure declarative programming languages due to their inability to express destructive assignment in constant time.

When the UF algorithm is extended to deal with function terms (rational trees), the algorithm can be used for optimal complexity *unification* [MF07]. Last but not least, a generalization of Union-Find that yields novel incremental algorithms for simple Boolean and linear equations [Frü06]. See chapter 10 in [Frü09] for an overview of Union-Find in CHR.

Parallelizing Basic Union-Find. For space reasons, we only discuss the basic UF algorithm here, not the optimized one. In CHR, the *data constraints* `root` and `->` represent the tree data structure. With the UF algorithm come several *operation constraints*: `find` returns the root of the tree in which a node is contained, `union` joins the trees of two nodes, `link` performs the actual join.

```
union    : union(A,B) <=> find(A,X), find(B,Y), link(X,Y).

findNode : A->B  \ find(A,X) <=> find(B,X).
findRoot : root(A) \ find(A,X) <=> found(A,X).

linkEq   : link(X,Y), found(A,X), found(A,Y) <=> true.
linkRoot : link(X,Y), found(A,X), found(B,Y), root(A) \ root(B) <=> B->A.
```

The second argument of the find operation `find` holds a fresh variable as identifier. When the root is found, it is recorded in the constraint `found`.

Standard CHR confluence analysis produces abstract states that reveal a deadlock: when we are about to apply the `linkRoot` rule, another link operation may remove one of the roots that we need for linking. From the non-confluent states we can derive an additional rule for `found` that mimics the rule `findNode`: the `found` constraint now keeps track of the updates of the tree so that its result argument is always a root.

```
foundUpdate : A->B \ found(A,X) <=> found(B,X).
```

Linking for disjoint node pairs can now run in parallel. While this seems an obvious result, this semi-automatic confluence-based approach yields a non-trivial parallel variant of the optimized UF algorithm with path compression. Correctness of the parallelisation is proven in both cases.

4 Parallel CHR with Transactions

Transactions alleviate the complexity of writing concurrent programs by offering entire computations to run atomically and in isolation. *Atomicity* means that a transaction either successfully commits or has to rollback (undo). In *optimistic concurrency control*, updates are logged and only committed at the end of an transaction when there are no update conflicts with other transactions. *Isolation* means that no intermediate update is observable by another transaction. The highest level of isolation is *serializability*, the major correctness criterion for concurrent transactions: for each parallel execution there is a sequential execution with the same result.

We now extend parallel CHR by transactions. Transactions will also be used for the implementation of parallel CHR in section 6 and for encoding of a transaction-basedconcurrency model in CHR in section 10.1.

4.1 Transactions in Parallel CHR

The paper [SS08] proposes CHRt as a conservative extension of CHR with atomic transactions. An atomic transaction is denoted as a meta-constraint `atomic(C)` where `C` is a conjunction of CHR constraints. Atomic transactions may appear in goals.

Example 4.1. Consider these CHR rules for updating a bank account:

```
balance(Acc,Bal), deposit(Acc,Amt)  <=> balance(Acc,Bal+Amt).
balance(Acc,Bal), withdraw(Acc,Amt) <=> Bal>Amt | balance(Acc,Bal-Amt).

transfer(Acc1,Acc2,Amt) <=> withdraw(Acc1,Amt), deposit(Acc2,Amt).
```

The `balance` constraint is a *data constraint*, and the `deposit` and `withdraw` constraints are *operation constraints*. The guard ensures that withdrawal is only possible if the amount in the account is sufficient. The transfer constraint rule combines deposit and withdrawal among two accounts.

Now assume a transfer between two accounts:

```
balance(acc1,500), balance(acc2,0), transfer(acc1,acc2,1000)
```

We can execute the deposit, but we cannot execute the withdrawal due to insufficient funds. The transaction gets *stuck*. This is clearly not the desired behavior of a transfer.

In CHRt, we can introduce a transaction to avoid this problem. The `transfer` constraint in the goal is wrapped by the meta-constraint `atomic`.

```
balance(acc1,500), balance(acc2,0), atomic(transfer(acc1,acc2,1000))
```

Now the incomplete transaction will be rolled back, no money will be transferred.

4.2 Abstract Syntax and Semantics of CHRt

We assume Ground CHR. We distinguish between data constraints and operation constraints. The head of a CHRt rule must contain exactly one operation constraint.

It requires one more transition for transactions. The $(Atomic)$ transition executes any number of atomic transactions in parallel.

$$\text{(Atomic)} \qquad \frac{T \wedge S_1 \wedge C_1 \mapsto^* T \wedge S_1', \ldots T \wedge S_n \wedge C_n \mapsto^* T \wedge S_n'}{T \wedge S_1 \wedge \ldots S_n \wedge atomic(C_1) \wedge \ldots atomic(C_n) \Rightarrow T \wedge S_1' \wedge \ldots S_n'}$$

T, S_i, S_i' must be data constraints. So there are only data constraints in the target state. The parallel step considers the separate evaluation of each C_i in isolation. The transactions only share the common data constraints T, which serves as a context. Note that each transaction may perform arbitrary many computation steps. Each transaction is fully executed until there are no operation constraints, it does not not get stuck.

4.3 Monotonicity, Soundness and Serializability

For CHRt programs, the following properties are proven to hold.

Serializability For each (Atomic) transition with n concurrent transactions, there is a corresponding computation of n consecutive sequential (Atomic) transitions each with only one transaction.

Soundness For any computation in CHRt, there is a corresponding computation in CHR where the atomic wrappers are dropped.

Monotonicity Although not proven in the paper, it follows from Soundness and the context T of the (Atomic) transition.

4.4 Encoding Transactions in Standard CHR

We want to execute CHRt in standard parallel CHR, i.e without the (Atomic) transition. The straightforward way is to execute atomic transactions only sequentially. Thus, we trivially guarantee the atomic and isolated execution of transactions. We identify two special cases where we can erase the atomic wrappers and still allow for parallel execution: bounded and for confluent transactions.

Bounded Transactions. A *bounded transaction* is one that performs a finite, statically known number of transitions. We eliminate a bounded transaction `atomic(G)` from a program by adding a rule to the program of the form `atomic(G) <=> G`. Then we unfold the rule [Frü05b, GMTW13] until no more operation constraints appear in its body. Since the transaction is bounded, unfolding will eventually stop.

In the running example, we can replace the atomic transfer rule (since it is bounded) by the following rule.

```
balance(Acc1,Amt1),balance(Acc2,Amt2),atomic(transfer(Acc1,Acc2,Amt)) <=>
             Amt1>Amt | balance(Acc1,Amt1-Amt), balance(Acc2,Amt2+Amt).
```

The rule head expresses the fact that an atomic transfer requires exclusive access to both accounts involved.

Confluent Transactions. The paper proves that if a CHRt program is confluent when we ignore atomic wrappers, then it can be executed in standard parallel CHR provided the initial goal never gets stuck. *Confluence* then guarantees that isolation is not violated.

Consider the example of the stuck transaction that attempts to overdraw an account. The withdraw rule can be fixed if we drop its guard (and hence allow negative balances): Any two consecutive transfers commute now. Regardless of the order they are performed in, they yield the same final result (even if the intermediate results differ). Hence, we can safely erase the atomic wrappers.

5 A More Refined Parallel Semantics

A *refined* semantics for parallel CHR is developed and implemented in [LS07, SL08, LS09, Lam11a]. This semantics can be seen as a refinement of the parallel abstract semantics given before. As before, the semantics is given in two parts, the sequential transitions and the parallel transitions.

5.1 Syntax for Refined Parallel CHR

Built-In Constraint	e		
CHR Constraint	$c ::= p(\overline{t})$	Identified (CHR) Constraint	$nc ::= c\#i$
Goal Constraint	$g ::= c \mid e \mid nc$	Goal (Store)	$G ::= \biguplus g$
Store Constraint	$sc ::= e \mid nc$	(Constraint) Store	$Sn ::= \bigcup sc$
State	$\sigma ::= \langle G, Sn \rangle$	Matched Constraints	$\delta ::= Sn \setminus Sn$

Figure 1: Refined Parallel CHR Syntax

Figure 1 describes the syntax for the refined semantics. The notation $\overline{a}$ denotes a sequence of a's. We only consider built-in constraints that are syntactic or arithmetic equations. To distinguish multiple occurrences (copies, duplicates) of CHR constraints, they are extended by a unique identifier. We call $c\#i$ an *identified constraint*. Unlike in the abstract semantics, conjunctions are modeled as (multi)sets and a state is now a pair: we distinguish between a goal (store) (a multiset of constraints) and the (constraint) store (a set of built-in and identified CHR constraints). We also introduce matched constraints that we will need as an annotation to transitions.

5.2 Sequential Refined CHR Semantics

The sequential part of the semantics in Figure 2 is a generalization of the refined CHR semantics of [DSGH04]. The semantics assumes generalized simpagation rules that are not propagation rules.

Constraints from the goal are executed one by one. Such a constraint is called *active*. It tries to apply rules to itself. To try a rule, the active constraint is matched against a head constraint of the rule. The remaining head constraints are matched with *partner constraints* from the constraint store. If there is such a complete matching and if the guard is satisfied under this matching, then the rule applies (fires). The constraints matching the removed constraints of the head are deleted and the body of the rule is added to the state. Because of the role of the active constraint, we call the semantics *goal-based semantics*.

Transitions. A transition $\sigma \overset{\delta}{\longmapsto} \sigma'$ maps the CHR state σ to σ' involving the CHR constraint goals in δ. The transition annotation δ holds the constraints that where matched with the rule head. It will be needed in the parallel part of the semantics.

The first transition **(Solve+Wake)** moves an equation built-in constraint e into the store and wakes up (reactivates) identified constraints in the store which could participate in a rule application. The function $WakeUp(e, Sn)$ computes a conservative approximation of the reactivated constraints, where m.g.u. denotes the most general unifier induced by a set of equations.

In transition **(Activate)**, a CHR constraint goal becomes active by annotating it with a fresh unique identifier and adding it to the store.

Rules are applied in transitions **(Apply-Remove)** and **(Apply-Keep)**. They are analogous, but distinguish if the active constraint $c\#i$ is kept or removed. In both cases, we seek for the missing partner constraints in the store, producing a *matching substitution* ϕ in case of success. The guard t must be logically entailed by the built-in constraints in the store under the substitution ϕ. Then we apply the rule instance of r by removing the matching constraints H_S and adding the rule body instance $\phi(B)$ to the goal. We also record the matched identified constraints H_S and H_P in the transition annotation. In transition **(Apply-Remove)**, the matching constraints H_S include $c\#i$. Since $c\#i$ is removed, we drop it from both the goal and the store. In transition **(Apply-Keep)**, $c\#i$ remains and so can possibly fire further rules.

Finally, in transition **(Suspend)**, we remove an active identified constraint from the goal if no (more) rules apply to the constraint. Note that the constraint is still (suspended) in the store and may be woken later on.

(Solve+Wake)
$$\frac{W = WakeUp(e, Sn)}{\langle \{e\} \uplus G \mid Sn \rangle \overset{W\setminus\{\}}{\rightarrowtail} \langle W \uplus G \mid \{e\} \cup Sn \rangle}$$

(Activate)
$$\frac{i \text{ is a fresh identifier}}{\langle \{c\} \uplus G \mid Sn \rangle \overset{\{\}\setminus\{\}}{\rightarrowtail} \langle \{c\#i\} \uplus G \mid \{c\#i\} \cup Sn \rangle}$$

(Apply-Remove)
$$\frac{\begin{array}{c}\text{Variant of } (r \ : \ H'_P \backslash H'_S <=> t \mid B') \in \mathscr{P} \text{ such that} \\ \exists \phi \quad Eqs(Sn) \models \phi(t) \quad \phi(H'_P) = DropIds(H_P) \\ \phi(H'_S) = \phi(\{c\} \uplus DropIds(H_S)) \quad \delta = H_P \backslash \{c\#i\} \cup H_S\end{array}}{\begin{array}{c}\langle \{c\#i\} \uplus G \mid \{c\#i\} \cup H_P \cup H_S \cup Sn \rangle \\ \overset{\delta}{\rightarrowtail} \quad \langle \phi(B') \uplus G \mid H_P \cup Sn \rangle\end{array}}$$

(Apply-Keep)
$$\frac{\begin{array}{c}\text{Variant of } (r \ : \ H'_P \backslash H'_S <=> t \mid B') \in \mathscr{P} \text{ such that} \\ \exists \phi \quad Eqs(Sn) \models \phi(t) \quad \phi(H'_S) = DropIds(H_S) \\ \phi(H'_P) = \phi(\{c\} \uplus DropIds(H_P)) \quad \delta = \{c\#i\} \cup H_P \backslash H_S\end{array}}{\begin{array}{c}\langle \{c\#i\} \uplus G \mid \{c\#i\} \cup H_P \cup H_S \cup Sn \rangle \\ \overset{\delta}{\rightarrowtail} \quad \langle \phi(B') \uplus \{c\#i\} \uplus G \mid \{c\#i\} \cup H_P \cup Sn \rangle\end{array}}$$

(Suspend)
$$\frac{\textbf{(Apply-Remove)} \text{ and } \textbf{(Apply-Keep)} \text{ do not apply to } c\#i \text{ in } Sn}{\langle \{c\#i\} \uplus G \mid Sn \rangle \overset{\{\}\setminus\{\}}{\rightarrowtail} \langle G \mid Sn \rangle}$$

where
$$\begin{array}{lcl} Eqs(S) & = & \{e \mid e \in Sn, e \text{ is an equation}\} \\ DropIds(Sn) & = & \{c \mid c\#i \in Sn\} \uplus \{e \mid e \in Sn\} \\ WakeUp(e, Sn) & = & \{c\#i \mid c\#i \in Sn \wedge \phi \text{ m.g.u. of } Eqs(Sn) \wedge \\ & & \theta \text{ m.g.u. of } Eqs(Sn \cup \{e\}) \wedge \phi(c) \neq \theta(c)\} \end{array}$$

Figure 2: Parallel CHR Semantics (Sequential Part $\overset{\delta}{\rightarrowtail}$)

5.3 Extension to Parallel Refined CHR Semantics

Figure 3 presents the parallel part of the refined operational semantics. It is a refined of the parallel transition for the abstract semantics. We allow for multiple goal stores while the constraint store is shared among the parallel computations.

In the **(Lift)** transition, we turn a sequential computation into a parallel computation. Transition **(Parallel-Goal)** parallelizes two parallel computations operating on the same shared store, if their matched constraints δ_1 and δ_2 do not have an overlap that involves removed constraints. They may overlap in the kept constraints. This makes sure that parallel computations remove distinct constraints in the store. The identifiers of constraints make sure that we can remove multiple but different copies of the same constraint. The matched constraints δ_1 and δ_2 are composed by the union of the kept and removed components, respectively, forming δ. Note that a context G is added to the goals in the resulting parallel transition, implying monotonicity.

5.4 Monotonicity, Soundness and Serializability

Monotonicity holds for the goal store, but not for the constraint store. In an enlarged constraint store, the **(Suspend)**

(Lift)

$$\frac{\langle G \mid Sn\rangle \overset{\delta}{\rightarrowtail} \langle G' \mid Sn'\rangle}{\langle G \mid Sn\rangle \overset{\delta}{\rightarrowtail}_{||} \langle G' \mid Sn'\rangle}$$

(Parallel-Goal)

$$\frac{\begin{array}{c}\langle G_1 \mid H_{S1} \uplus H_{S2} \uplus Sn\rangle \overset{\delta_1}{\rightarrowtail}_{||} \langle G_1' \mid H_{S2} \uplus Sn\rangle \\ \langle G_2 \mid H_{S1} \uplus H_{S2} \uplus Sn\rangle \overset{\delta_2}{\rightarrowtail}_{||} \langle G_2' \mid H_{S1} \uplus Sn\rangle \\ \delta_1 = H_{P1} \backslash H_{S1} \quad \delta_2 = H_{P2} \backslash H_{S2} \\ H_{P1} \subseteq Sn \quad H_{P2} \subseteq Sn \quad \delta = H_{P1} \uplus H_{P2} \backslash H_{S1} \uplus H_{S2} \\ H_{S1} \cap (H_{P2} \uplus H_{S2}) = \{\} \quad H_{S2} \cap (H_{P1} \uplus H_{S1}) = \{\}\end{array}}{\langle G_1 \uplus G_2 \uplus G \mid H_{S1} \uplus H_{S2} \uplus Sn\rangle \overset{\delta}{\rightarrowtail}_{||} \langle G_1' \uplus G_2' \uplus G \mid Sn\rangle}$$

Figure 3: Parallel CHR Semantics (Parallel Part $\overset{\delta}{\rightarrowtail}_{||}$)

transition may not be possible anymore, because a new rule becomes applicable to the active constraint. The monotonicity is still sufficient though, because in the semantics, the constraint store is only populated via the goal store.

Serializability holds: Any parallel computation can be simulated by a sequence of sequential computations in the refined semantics. Furthermore, soundness holds: any parallel computation has a correspondence in a suitable variant of the sequential abstract semantics. For the upcoming theorem, let us note that an *initial state* is of the form $\langle G, \{\}\rangle$, a *final state* is of the form $\langle \{\}, Sn\rangle$. Given a computation $\langle G \mid \{\}\rangle \rightarrowtail_{||}^{*} \langle G' \mid Sn\rangle$, the state $\langle G' \mid Sn\rangle$ is called a *reachable state*.

Theorem 5.1 (Soundness)**.** For any reachable state $\langle G \mid Sn\rangle$,

$$\begin{array}{ll}\text{if} & \langle G \mid Sn\rangle \rightarrowtail_{||}^{*} \langle G' \mid Sn'\rangle \\ \text{then} & (NoIds(G) \uplus DropIds(Sn)) \mapsto^{*} (NoIds(G') \uplus DropIds(Sn'))\end{array}$$

where $\mapsto^{*}$ denotes transitions in the sequential abstract semantics and where $NoIds = \{c \mid c \in G, c \text{ is a CHR constraint}\} \uplus \{e \mid e \in G, e \text{ is an equation}\}$.

6 Parallel CHR Implementation in Haskell

The parallel refined semantics from the previous section 5 has been implemented in the lazy functional programming language Haskell [SL07, LS07, SL08, LS09, Lam11a]. We use the Glasgow Haskell Compiler for implementing parallel Ground CHR. In principle, the system can be reimplemented in mainstream languages such as C and Java. The implementation is available for download at `https://code.google.com/archive/p/parallel-chr/`.

6.1 Implementation Principles

Our implementation follows the principles of standard sequential implementations of CHR where possible [HGSD05, VW10] including their optimizations, such as constraint indexing (hashing), optimal join ordering and early guard scheduling.

Parallel goal execution must not remove constraints in overlaps that participate in several rule head matchings. We discuss two approaches to implement this kind of parallel rule-head matching, locking and transactions, before we settle for a hybrid approach.

Fine-grained Lock-based Parallel Matching. We restrict the access to each constraint in the shared store with a lock. An active constraint will first try to lock its matching removed partner constraints. Kept constraints can be used

by several rules simultaneously, but must not be locked for removal. Locking fails if any constraint in the complete rule head matching is already locked by another active constraint. If locking fails, we release the locks and try to redo the rule application. Otherwise, the rule is applied. No unlocking is necessary since locked constraints are removed. This locking mechanism can avoid deadlocks and cyclic behavior using standard techniques for these problems such as timestamps.

Software Transactional Memory (STM). Instead of using locks directly, we can use the STM transactions provided in Haskell. The principles of transaction have been introduced in section 4. If there are update conflicts among transactions, the STM protocol will randomly commit one of the atomic transactions and rollback the others. The disadvantage of STM is that unnecessary rollbacks can happen. We will meet STM again in section 10.1, when it is specified in CHR.

Hybrid STM-based Locking Scheme. In the implementation, we use both Software Transactional Memory (STM) and traditional shared memory access locking techniques. The search for matching partner constraints is performed outside STM to avoid unnecessary rollbacks. When a complete rule head matching is found, we perform an STM procedure that we call *atomic rule-head verification (ARV)*. It checks that all the constraints are still available and marks the constraints to be removed as deleted. These deleted constraints will be physically delinked from the constraint store, either immediately or later by an amortized removal procedure. Both can be implemented with standard concurrency primitives (such as compare-and-swap and locks).

Thread Pool. The naive way to implement a parallel CHR system is to spawn an active thread for each goal constraint in a state. Each thread tries to find its partner constraints. However, the thread and its later partner constraints would then compete for the same rule application. Moreover, the number of threads would be unbounded, as the number of constraints in a state is unbounded. Our implementation uses a bounded number of active threads. A *thread pool* maintains threads waiting for tasks to be allocated for parallel execution.

6.2 Experimental Results

Experiments were performed on an Intel Core quad core processor with hyper-threading technology (that effectively allows it to run 8 parallel threads). We measure the relative performance of executing with 1, 2, 4, 8 and an unbounded number of threads against a sequential CHR implementation in Haskell. The table in Figure 4 gives some exemplary results with optimizations: each goal thread searches store constraints in a unique order to avoid matching conflicts and a special goal ordering for Merge Sort and Gcd is used (explained below).

Number of Threads	1	2	4	8	Unbounded
Merge Sort	121%	94%	70%	52%	>200%
Gcd	109%	37%	18%	12%	123%
Parallel Union-Find	125%	82%	52%	32%	>200%
Blockworld	123%	77%	54%	39%	>200%
Dining Philosophers	119%	74%	49%	41%	>200%
Prime	115%	73%	46%	30%	155%
Fibonacci	125%	85%	59%	39%	>200%
Turing Machine	111%	63%	78%	70%	>200%

Figure 4: Experimental results, with optimal configuration (on 8 threaded Intel processor)

There are several general observations to be made with regard to the number of threads. Executing with one goal thread is clearly inferior to the sequential implementation because of the wasted overhead of parallel execution. Executions with 2, 4 and 8 goal threads show a consistent parallel speed-up, with exception of the Turing Machine. It is inherently single-threaded. Interestingly, we still obtain improvements from parallel execution of administrative procedures. Unbounded thread pooling is always slower than the sequential implementation. Furthermore we observed a super-linear speed-up for the Gcd example with a *queue-based goal ordering* instead of the usual stack-based ordering in the goal store. In merge sort, we stack `->` constraints and queue just `=>` for optimal performance. Last but not least,

experiments also confirmed that there is a speed-up when a multi-core processor instead of a single-core processor is used.

7 Massively-Parallel Set-Based CHR Semantics

A CHR semantics is *set-based* if conjunctions of constraints are considered as set instead of multiset. In [RF10], we present a parallel execution strategy for set-based CHR. The use of sets instead of multisets has a dramatic impact: it allows for multiple removals of constraints. This means that overlaps can be removed several times. We show that the resulting refined semantics is sound with respect to sequential execution if the program is deletion-acyclic (i.e., when there is no mutual removal of constraints). CHRmp programs for the computation of minimum, prime numbers, and sorting can run in constant time, given enough processors. Similarly, we describe a program for SAT solving in linear time.

7.1 Massively-Parallel Set-Based Semantics

As in the parallel abstract semantics, there are no restrictions on the syntax of CHR. Reconsider the essential (Parallel) transition of the abstract CHR semantics. Keep in mind that conjunctions of constraints are now interpreted as sets of constraints.

$$\text{(Parallel)} \quad \frac{A \wedge E \mapsto C \wedge E \qquad B \wedge E \mapsto D \wedge E}{A \wedge B \wedge E \mapsto C \wedge D \wedge E}$$

Consider the program

```
a <=> b,c.          a <=> b,d.
```

Then the following transition for the goal $\mathtt{a} \wedge \mathtt{e}$ is possible in the set-based interpretation:

$$\frac{\mathtt{a} \wedge \mathtt{e} \mapsto \mathtt{b} \wedge \mathtt{c} \wedge \mathtt{e} \qquad \mathtt{a} \wedge \mathtt{e} \mapsto \mathtt{b} \wedge \mathtt{d} \wedge \mathtt{e}}{\mathtt{a} \wedge \mathtt{e} \mapsto \mathtt{b} \wedge \mathtt{c} \wedge \mathtt{d} \wedge \mathtt{e}}$$

This means that a is removed twice and b is only produced once.

When we generalize this observation, we see that overlaps between rule matchings can be removed arbitrary many times, leading to a kind of massive parallelism.

We will refine this set-based semantics now. We assume CHR without propagation rules. In the body of a rule, we distinguish between CHR constraints B_c and built-in constraints B_b, and write B_c, B_b. A CHRmp state is of the form $\langle \mathbb{G}; \mathbb{B} \rangle$, where the goal (store) $\mathbb{G}$ is a *set* (not multiset) of constraints and the (built-in) constraint store $\mathbb{B}$ is a conjunction of built-in constraints. We adapt the state equivalence $\equiv$ in the obvious way to CHRmp states.

Definition 7.1 (Massively Parallel Transition)**.** Given a CHRmp state $\sigma = \langle \mathbb{G}; \mathbb{B} \rangle$. Let $\mathscr{R}$ be the smallest set such that for each rule variant $r : H_1 \backslash H_2 \Leftrightarrow G \mid B_c, B_b$ where $\sigma \equiv \langle H_1 \cup H_2 \cup \mathbb{G}'; G \wedge \mathbb{B}' \rangle$ it holds that $(H_1, H_2, B_c, B_b, \mathbb{B}') \in \mathscr{R}$. We then define for any non-empty subset $R \subseteq \mathscr{R}$:

- the set of removed constraints $D = \{c \mid \exists (_, H_2, _, _, \mathbb{B}') \in R, c \in \mathbb{G} : H_2 \wedge \mathbb{B}' \rightarrow c\}$
- the set of added constraints $A = \{c \mid \exists (_, _, B_c, _, _) \in R : c \in B_c\}$
- the conjunction of added built-in constraints $B = \bigwedge_{(_,_,_,B_b,\mathbb{B}') \in R} \mathbb{B}' \wedge B_b$

A *massively parallel transition (step)* of $\sigma = \langle \mathbb{G}; \mathbb{B} \rangle$ using $\mathscr{R}$ is then defined as:

$$\text{(Massive-Apply)} \quad \langle \mathbb{G}; \mathbb{B} \rangle \twoheadrightarrow^R \langle (\mathbb{G} \setminus D) \cup A; \mathbb{B} \wedge B \rangle$$

If the specific set R is not of importance we write $\twoheadrightarrow$ instead of $\twoheadrightarrow^R$.

In the extreme case, $R = \mathscr{R}$, so all possible rule applications are performed simultaneously. We call this *exhaustive parallelism*. With such an execution strategy, any CHRmp program is trivially *confluent*, because there are no conflicting rule applications. On the other hand, if R is a singleton set, only one rule is applied and we are back to sequential CHR.

Example 7.1. Reconsider the CHR program for computing prime numbers. Consider the state

$$\sigma = \langle \{prime(2), prime(3), prime(4), prime(5), prime(X)\}; X{=}6 \rangle.$$

There are three possible rule applications, removing the non-prime numbers 4 and twice 6:

$$\mathscr{R} = \left\{ \begin{array}{l} (\{prime(N_1)\}, \{prime(M_1)\}, \emptyset, \top, X{=}6 \wedge N_1{=}2 \wedge M_1{=}4), \\ (\{prime(N_2)\}, \{prime(M_2)\}, \emptyset, \top, X{=}6 \wedge N_2{=}2 \wedge M_2{=}6), \\ (\{prime(N_3)\}, \{prime(M_3)\}, \emptyset, \top, X{=}6 \wedge N_3{=}3 \wedge M_3{=}6) \end{array} \right\}$$

We can now perform all three possible rule applications exhaustively parallel, i.e. $R = \mathscr{R}$, resulting in the following sets:

$$D = \{prime(4), prime(X)\}, \;\; A = \emptyset,$$
$$B = (X{=}6 \wedge N_1{=}2 \wedge M_1{=}4) \wedge (X{=}6 \wedge N_2{=}2 \wedge M_2{=}6) \wedge (X{=}6 \wedge N_3{=}3 \wedge M_3{=}6)$$

This leads to the parallel transition:

$$\sigma \twoheadrightarrow^{\mathscr{R}} \quad \langle \{prime(2), prime(3), prime(5)\}; X{=}6 \wedge B \rangle$$

Hence, a single parallel step is sufficient to find all prime numbers.

7.2 Example Programs under Exhaustive Parallelism

We examine different algorithms written in CHR and the effect of executing these programs in CHRmp, in particular with exhaustive parallelism to achieve maximum speed-up.

Filter Programs. Programs that only consist of rules whose body is `true` can be understood as filtering constraints. They can obviously be executed in constant time with exhaustive parallelism, given enough processors. The prime program falls into this category. The experiments with the prime program using massive parallelism [TORF12] show an run-time improvement of about an order of magnitude over strong parallelism. Another example is the computation of the minimum.

Sorting. The `msort` rule of merge sort leads to a linear number of exhaustively parallel steps. We can significantly improve on that by the following variant of the program. We first compute all ordered pairs (this also happens with `msort`), but then filter out all pairs that do not belong to the final result. This can be accomplished by introducing a second type of arc pair constraints to distinguish the two steps and these two rules:

```
allpairs   : A->B \ A->C <=> A<B, B<C | B>>C.
filtersort : A>>B \ A>>C <=> A<B, B<C | true.
```

Note that the second rule can be seen as a minimum computation among all numbers greater than `A`. In this way, the direct successor of `A` in the sorted chain can be found. This program can sort in two exhaustively parallel steps, thus takes constant time.

SAT Solving. The SAT formula is given as a tree of its subexpressions. The tree nodes are of the form `eq(Id,B)`, where `Id` is a node identifier and `B` is either a Boolean variable written `v(X)` or a Boolean operation (`neg, and, or`) applied to identifiers. Additionally, a `f(L,[])` constraint is required in the initial state, where `L` is a list of all n variables in the SAT formula.

```
generate : f([X|Xs], A) <=> f(Xs,[true(X)|A]), f(Xs,[false(X)|A]).
assign   : f([],A) \ eq(T,v(X)) <=> true(X) in A | sat(T,A,true).
assign   : f([],A) \ eq(T,v(X)) <=> false(X) in A | sat(T,A,false).

sat(T1,A,S) \ eq(T,neg(T1)) <=> sat(T,A, neg S).
sat(T1,A,S1), sat(T2,A,S2) \ eq(T,and(T1,T2)) <=> sat(T,A, S1 and S2).
sat(T1,A,S1), sat(T2,A,S2) \ eq(T,or(T1,T2)) <=> sat(T,A, S1 or S2).
```

The `generate` rule generates, in n parallel steps, 2^n `f` constraints representing all possible truth assignments to variables as a list in its second argument. In the next parallel step (using the `assign` rules) all n Boolean variables in the given formula are assigned truth values for each assignment, represented by `sat` constraints.

The remaining three rules determine the truth values of all sub-expressions of the formula bottom-up. In each parallel step the truth values of sub-expressions at a certain height of the tree are concurrently computed for all possible assignments of variables. Therefore, the number of parallel steps in this phase is bound by the depth of the formula.

A formula is in 3-DNF normal form if it is in disjunctive normal form (a disjunction of conjunctions of literals) and each clause contains at most 3 literals. Because of its bounded depth, a SAT problem given in 3-DNF normal form with n variables can be solved in linear time in n with this program under exhaustive parallelism, independent of the size of the formula.

7.3 Soundness and Deletion-Acyclicity

Soundness of CHRmp is not always possible as the following example shows.

Example 7.2. Consider the following rule that removes one of two differing constraints,

```
c(N) \ c(M) <=> N=\=M | true.
```

and the goal `c(1), c(2)`. There are two competing rule instances for application: one matches the two constraints in the given order, the other in reversed order. So if we apply both rules simultaneously under exhaustive parallelism, both constraints will be (incorrectly) removed.

In general, computations that allow for mutual removal of constraints are not sound in CHRmp. Soundness therefore requires that the programs are deletion-acyclic. A *deletion dependency pair* (c,d) means the constraint c is required to remove constraint d in some rule.

Definition 7.2 (Deletion Dependency, Deletion-Acyclic)**.** Let $\sigma = \langle \mathbb{G}; \mathbb{B} \rangle \twoheadrightarrow^R \tau$, then the *deletion dependency* $\mathscr{D}(\sigma, R)$ is a binary relation such that $(c,d) \in \mathscr{D}(\sigma,R)$ if and only if $c,d \in \mathbb{G}$ and there exist $(H_1,H_2,B_c,B_b,\mathbb{B}') \in \mathscr{R}$ and $c' \in H_1, d' \in H_2$ such that $c' \wedge \mathbb{B}' \to c$ and $d' \wedge \mathbb{B}' \to d$.

A CHRmp program $\mathscr{P}$ is *deletion-acyclic* if and only if for all σ such that $\sigma \twoheadrightarrow^{\mathscr{R}} \tau$ the transitive closure $\mathscr{D}(\sigma,\mathscr{R})^+$ is irreflexive.

In a deletion-acyclic program, we can simulate the CHRmp computation steps by a sequence of sequential rule applications in multiset semantics, provided we initially have enough copies of the user-defined constraints and can remove them when needed. The latter is accomplished by rules of the form

```
set-rule: c(X1,...Xn) \ c(X1,...Xn) <=> true.
```

for each CHR constraint `c/n` in the given program. We call these *set-rules*. Let $\rightarrowtail$ be a sequential transition in a suitable variant of the usual multiset CHR semantics.

Theorem 7.1 (Soundness)**.** Let $\mathscr{P}$ be a deletion-acyclic CHRmp program and $\mathscr{P}'$ be the CHR program $\mathscr{P}$ extended with set-rules. If $\sigma = \langle \mathbb{G}; \mathbb{B} \rangle \twoheadrightarrow_{\mathscr{P}} \tau$, then there exists a multiset $\mathbb{G}'$ with $c \in \mathbb{G}' \Leftrightarrow c \in \mathbb{G}$ such that $\sigma' = \langle \mathbb{G}'; \mathbb{B} \rangle \rightarrowtail^*_{\mathscr{P}'} \tau'$, where $c \in \tau' \Leftrightarrow c \in \tau$.

Example 7.3. Consider the initial goal `a` and the program

```
a <=> b,c.          a <=> b,d.          b,c,d <=> true.
```

Exhaustive parallelism leads to the set-based computation
`a` $\twoheadrightarrow$ `b,c,d` $\twoheadrightarrow$ true.

The sequential correspondence in the multiset CHR program extended with set-rules is

`a,a` $\rightarrowtail$ `b,b,c,d` $\rightarrowtail$ `b,c,d` $\rightarrowtail$ true.

The example can also be used to show that *Serializability* in general does not hold for massively-parallel set-based CHR. There is not sequential computation in CHRmpthat can simulate the exhaustively parallel computation, since the first rule application will remove `a`, so either `b,c` or `b,d` can be produced sequentially, but not their union. Similarly, monotonicity does not hold.

8 Parallel Hardware Implementations

The work reported in [TORF12] investigates the compilation of CHR to specialized hardware. We translate the CHR code into the low-level hardware description language VHDL, which in turn creates the necessary hardware using Field Programmable Gate Array (FPGA) technology. FPGA is a hardware consisting of programmable multiple arrays of logic gates. We also implement a hybrid CHR system consisting of a software component running a CHR system for sequential execution, coupled with hardware for parallel execution of dedicated rules in the program. The resulting hardware system is typically an order of magnitude faster than the fastest software implementation of CHR.

8.1 Basic Compilation of CHR to Procedural Languages

As preliminaries, we give the basic implementation scheme for CHR in procedural languages like C and Java, but also VHDL. This scheme applies throughout this section. A CHR rule can be translated into a procedure using the following simple scheme:

```
procedure(rule_kept_head_constraints, removed_head_constraints) {
    if (head constraints not marked removed  &&  head matching  &&  guard check)
    then  execute built-in constraints;  add body CHR constraints;
}
```

The parameter list references the head constraints to be matched to the rule. In the procedure, we first check that the constraints have not been marked as removed. Then head matching is explicitly performed and then the guard is checked. If all successful, one executes the built-in constraints and then adds the body CHR constraints. Added constraints may overwrite removed head constraints for efficiency. Constraints that are removed and not overwritten are marked as deleted. Such a rule procedure is executed on every possible combination of constraints from the store.

8.2 Compiling CHR to Parallel FPGA Hardware via VHDL

CHR Fragment with Non-Increasing Rules. We assume Ground CHR. Since the hardware resources can only be allocated at compile time, we need to know the largest number of constraints that can occur in the constraint store during the computation. In *non-increasing rules*, the number of body CHR constraints added is not greater than the number of head constraints removed. Thus the number of constraints in the initial goal provides an upper bound on the number of constraints during the computation. Hence we only allow for non-increasing simpagation rules.

CHR Compilation Hardware Components. A Program Hardware Block (PHB) is a collection of Rule Hardware Blocks (RHBs), each corresponding to a rule of the CHR program. A Combinatorial Switch (CS) assigns the constraints to the PHBs. In more detail:

Rule Hardware Block (RHB) In VHDL the rule is translated into a single clocked process following the transformation scheme described above. Here, the parameters are input signals for each argument of the head constraints. Each signal is associated with a validity signal to indicate if the associated constraint has been removed. A concrete example is given below.

Program Hardware Block (PHB) The PHB makes sure that the RHBs keep applying themselves until the result remains unchanged for two consecutive clock cycles. Each rule is executed by one or more parallel processes that fire synchronously every clock cycle. The initial goal is directly placed in the constraint store from which several instances of the PHB concurrently retrieve the constraints.

Combinatorial Switch (CS) The CS sorts, partitions and assigns the constraints to the PHBs, ensuring that the entire constraint store gets exposed to the rule firing hardware. It acts as a synchronization barrier, allowing the faster PHBs to wait for the slower ones, then communicating the results between the blocks. It also reassigns the input signals to make sure that all constraint combinations have been exposed to the rule head matching.

Strong Parallelism with Overlap. For a given kept constraint, multiple RHBs are used to try rules with all possible partner constraints. For the case of simpagation rules with one kept and one removed constraint, we introduce a hardware block that consists of a circular shift register which contains all the initial goal constraints. The first register cell contains the kept constraint and it is connected to the first input of all the RHBs, the rest of the register cells contain the potential partner constraints and are each connected to the second input of one RHB. Every time the PHBs terminate their execution, the new added constraints replace the removed ones. They registers shift until a non-removed constraint is encountered.

Example 8.1. Consider the rule for the greatest common divisor:

```
r :  gcd(N) \ gcd(M) <=> M>=N | gcd(M-N).
```

In Figure 5 we give an excerpt of the VHDL code for the above rule. There are two processes executed in parallel, one for each matching order, that correspond to two RHBs called `r_1` and `r_2`. The input parameters `gcd1` and `gcd2` are byte signals holding the numbers. `valid1s` and `valid2s` are bit signals. They are set to `0` if the associated constraint is removed. The shared variable `flag` is a bit. It is used to control the application of the two processes.

```
r_1: process (..., gcd1s, gcd2s, valid1s, valid2s)
begin
  if ...      % checking and setting flags and parameters
    if (valid1s=1 and valid2s=1) then
      if gcd2s>=gcd1s then
        gcd2s <= gcd2s-gcd1s;
        flag := 1;
      else
        flag := 0;
      end if;
    end if;
  end if;
end process r_1;

r_2: process (..., gcd1s, gcd2s, valid1s, valid2s)
begin
  if ...      % checking and setting flags and parameters
    if (valid1s=1 and valid2s=1) then
      if flag=0 then
        if gcd1s>=gcd2s then
          gcd1s <= gcd1s-gcd2s;
        end if;
      end if;
    end if;
  end if;
end process r_2;
```

Figure 5: Excerpt of VHDL Code for GCD Rule

Massive Parallelism. The set-based semantics CHR^{mp} allows multiple simultaneous removals of the same constraint. Our implementation eliminates the conflicts in the constraint removals by allowing different rule instances to work concurrently on distinct copies of the constraints. We provide all possible combinations of constraints to distinct parallel PHB instances in a single step. So the same constraint will be fed to several PHBs. Valid constraints are collected. A constraint is valid if no PHB has removed it. This is realized in hardware by AND gates. The improvement due to massive parallelism is about an order of magnitude for goals with a low number of constraints and it decreases with higher numbers of constraints. This is due to reaching the physical bounds of the hardware.

Experimental Results. The FPGA implementations of CHR are at least one order of magnitude faster than fastest the software implementations of CHR. Strong parallelism improves the performance, and massive parallelism improves it even faster. In general, CHR-FPGA is slower but within the range of handcrafted VHDL code.

8.3 Parallel Execution on a Graphical Processing Unit GPU

Graphical Processing Units (GPUs) consist of hundreds of small cores, collectively operating to provide massive computation capabilities. Similar to the previous work on parallel CHR hardware, we transform *non-increasing Ground CHR rules* to C/C++, then to use a GPU to fire the rules on all combinations of constraints. As proof of concept, the scheme was encoded by hand for some typical CHR examples in [ZFG12].

Translation into C++ with CUDA. The constraint store is implemented as an array of fixed length consisting of the structures that represent CHR constraints. A CHR rule can be translated into a function in C/C++ using the previous translation scheme. The rule is executed on every possible combination of constraints using nested for-loops. Finally, the code is rewritten using the CUDA library. The outer for-loop is parallelized for the thread pools of the GPU.

Example 8.2. Recall the CHR rule for minimum:

```
minimum :  min(A) \ min(B) <=> A=<B | true.
```

The equivalent C/C++ function is shown below. No variable matching is needed in the rule.

```
void minimum(MIN &a, MIN &b) {
   if(!a.isRemoved && !b.isRemoved && a.value <= b.value)
       b.isRemoved = true;
}
```

The nested for-loop iterating over pairs of constraints is shown below:

```
void fire_minimum(MIN *min_store, int N) {
  for (int i = 0; i < N; i += 1)
    for (int j = 0; j < N; j += 1)
      if (i != j)
        minimum(min_store[i], min_store[j]);
}
```

9 Distribution in CHR

Before we introduce a full-fledged distributed refined semantics for CHR in detail, we set the stage by describing a distributed but sequential implementation of set-based CHR. This system is successfully employed in a verification system for concurrent software. Both semantics feature propagation rules, but they use different mechanisms to avoid their repeated re-application.

9.1 Distributed Set-Based Goal Stores in CHRd

CHRd [SSR07] is an implementation of a sequential *set-based refined* semantics for CHR with propagation rules. CHRd features a distributed constraint store.

Termination of Propagation Rules. There are basically two ways to avoid repeated application of propagation rules: Either they are not applied a second time to the same constraints or they do not add the same constraints a second time. Since we can remove constraints in CHR, usually the first option is chosen: we store the sequence of CHR constraint identifiers to which a propagation rule has been applied. It can be garbage-collected if one of the constraints is removed. This information is called a *propagation history*. CHRd replaces the check on the propagation history by an *occurrence check* on the constraint store. This can be justified by the set-based semantics.

Set-Based Refined Semantics. Our set-based semantics closely follows the standard refined semantics [DSGH04]. The essential differences are as follows:

- The propagation history is dropped from the states.
- There is an additional transition to ensure a set-based semantics. It removes a constraint from the goal store before its activation, if it is already in the constraint store.
- There are additional transitions to avoid immediate re-application of a propagation rule. In the first transition, all head matching substitutions where the active constraint is kept are computed at once and all corresponding rule instances are added to the goal store. These rule instances are called *conditional activation events*.
- When a conditional activation event is processed, it is checked if the matching head constraints are still in the constraint store. If not, a second transition removes the event from the goal store. Otherwise, a third transition applies the rule instance by adding its body constraints to the goal store.

The semantics does not model the distribution of the CHRd constraint store.

Our set-based semantics is not always equivalent to the standard refined semantics. In the semantics a propagation rule may fire again on an constraint that has been re-activated (woken). In the refined multiset semantics, it will not be fired again. So a CHR program may not terminate with the set-based semantics, but with the refined semantics.

Distributed Local Constraint Stores by Variable Indexing. Finding the partner constraints in head matching efficiently is crucial for the performance of a CHR system. If variables are shared among head constraints, we can use the corresponding arguments of the constraints for *indexing*. If the argument is an unbound variable at run-time, we store (a pointer to) the constraint as attribute of that variable. If the argument becomes bound (or even ground) at run-time, the constraint can be accessed from a hash table instead.

A conjunction of constraints is *direct-indexed* if all subsets of constraints share variables with the remaining constraints. A CHR program is direct-indexed if all its rule heads are direct-indexed. Clearly, head matching works best for direct-indexed programs. CHRd requires direct-indexed programs that only index on unbound variables. This permits the constraint store to be represented in a distributed fashion as a network of constraints on variables.

Any CHR program can be trivially translated to a direct-indexed program. We just have to add an argument to each CHR constraint that always contains the same shared variable. For example, the direct-index rule for minimum is:

```
min(X,N) \ min(X,M) <=> N=<M | true.
```

With the help of the new variable, we can distinguish between different minima. In general, this technique can be used to localize computations.

Implementation and Experimental Results. Our implementation CHRd has been integrated into XSB, a Prolog programming system with tabling. It can be obtained from `http://xsb.sourceforge.net`. CHRd performs significantly better on programs using tabling, and shows comparable results on non-tabled benchmarks. This indicates that constraint store occurrence checks can be done as efficiently as propagation history checks while avoiding the maintenance of a propagation history.

Verification of Multi-Threaded Applications. The paper [SSSD07] describes an approach for checking for deadlocks in multi-threaded applications based on the concurrency framework SynchroniZation Units MOdel (Szumo).

The framework associates each thread with a synchronization contract that governs how it must synchronize with other threads. At run-time, schedules are derived by negotiating contracts among threads.

The Szumo system includes a constraint solver written in CHRd encoding the synchronization semantics of thread negotiation. The verification system performs a reachability analysis: it constructs execution paths incrementally until either a deadlock is detected or further extending the path would violate a synchronization contract.

With Szumo, we analyzed an implementation of the dining philosophers problem, where no deadlock was found. We verified the in-order message delivery property of an n-place FIFO buffer. We also analyzed Fischers protocol, a mutual-exclusion protocol that is often used to benchmark real-time verification tools. There we employed CHRd to specify a solver for the clock constraints.

9.2 Distributed Parallel CHRe

The paper [LC13] introduces a decentralized distributed execution model consisting of an ensemble of computing entities, each with its own local constraint store and each capable of communicating with its neighbors: in CHRerules are executed at one location and can access the constraint stores of its immediate neighbors. We have developed a prototype implementation of CHRe in Python with MPI (Message Passing Interface) as a proof of concept and demonstrated its scalability in distributed execution. It is available at `https://github.com/sllam/msre-py`.

9.3 Syntax of CHRe

We assume Ground CHR. There are no built-in constraints in the body of a rule. CHRe introduces locations. All user-defined constraints in a program must be explicitly localized. A *location* l is a term (typically an unbound variable or constant) that annotates a CHR constraint c, written as $[l]c$. A location l is *directly connected* to a location l' if there is a constraint at location l that mentions location l' in one of its arguments.

We are interested in rules that can read data from up to n of their immediate neighbors, but writes to any number of neighbors. We therefore define *n-neighbor restricted rules*:

- The rule head refers to connected locations in a star topology. The primary location is in the center and directly connected to each n neighbors.
- If a variable in the head is shared between different locations, it also must occur in the primary location.
- Each constraint in the guard must only contain variables of the primary location and at most one other location of the rule head.

Example 9.1. This variant of the Floyd-Warshall algorithm computes all-pair shortest paths of a directed graph in a distributed manner.

```
base  : [X]arc(Y,D) ==> [X]path(Y,D).
elim  : [X]path(Y,D1) \ [X]path(Y,D2) <=> D1<D2 | true.
trans : [X]arc(Y,D1), [Y]path(Z,D2) ==> X\=Z | [X]path(Z,D1+D2).
```

We distinguish between arcs and paths. `[X]path(Y,D)` denotes a path of length `D` from `X` to `Y`. The rules `base` and `elim` are 0-neighbor restricted (local) rules because their left-hand sides involve constraints from exactly one location. Rule `trans` is a 1-neighbor restricted rule since its left-hand side involves `X` and a neighbor `Y`. We see that `X` is the primary location of this rule because it refers to location `Y` in an argument.

9.4 Abstract and Refined Semantics of CHRe

Abstract Distributed CHRe Semantics for n-Neighbor Restricted Rules. Each location has its own goal store. Based on the standard abstract CHR semantics, we introduce abstract *ensemble* states, which are sets of local stores G_k where G is a goal and k a unique location name. In the (Apply-n) transition, each of the locations in an n-neighbor rule provides a partial match in their stores. If the matchings can be combined and if the guard holds, we add the rule

body goals to their respective stores. We show *soundness* with respect to the standard CHR abstract semantics, where locations are encoded as an additional argument to each CHR constraint.

Refined Distributed CHRe Semantics for 0-Neighbor Restricted Rules. We extend the standard CHR refined semantics to support decentralized incremental multiset matching for 0-neighbor restricted rules.

Localized States. In CHRe, an *ensemble* Ω is a set of localized states. A *localized state* is a tuple $\langle \vec{\mathscr{U}}; \vec{\mathscr{G}}; \vec{\mathscr{S}}; \bar{\mathscr{H}} \rangle_k$, where

- the *Buffer* $\vec{\mathscr{U}}$ is a queue of CHR constraints that have been sent to a location,
- the *Goal Store (Execution Stack)* $\vec{\mathscr{G}}$ is a stack of the constraints to be executed,
- the *Constraint Store* $\vec{\mathscr{S}}$ is a set of identified constraints to be matched,
- the *Propagation History* $\bar{\mathscr{H}}$ is a set of sequences of identifiers of constraints that matched the head constraints of a rule.

To add a further level of refinement, an *active occurrenced CHR constraint* $c(\bar{x})\#i{:}j$ is an identified constraint that is only allowed to match with the j-th occurrence of the constraint predicate symbol c in the head of a rule of a given CHR program $\mathscr{P}$.

Figure not part of this draft

Figure 6: The Sequential Part of the Refined CHRe Semantics for 0-Neighbor Restricted Rule

Localized Sequential Transitions. Figure 6 shows the sequential transitions:

- The (Flush) transition step applies if the goal store is empty and the buffer is non-empty. It moves the buffer constraints into the goal store.
- The three transitions (Loc 1), (Loc 2) and (Loc 3) apply if the first goal in the goal store of location k is $[k']c$. They deliver constraint $[k']c$ to location k':
  - Step (Loc 1) applies if k' is distinct from k and there exists a location k'. It it strips the localization operator $[k]$ away and sends constraint c to the buffer of k'.
  - Step (Loc 2) applies if k' is the same as k. The localization operator $[k]$ is dropped.
  - Step (Loc 3) applies if there is no location k'. A new location k' is created with constraint c in the buffer.
- The remaining transitions apply to a location as to a state in the standard refined semantics. Buffers are ignored and remain unchanged. The transitions model the activation of a constraint, the application of rules to it, and its suspension if no more rule is applicable.
  - The (Act)ivate transition, a CHR constraint c becomes active (with first occurrence 1) and is also introduced as identified constraint into the constraint store.
  - The (Simp)lify transition applies a rule where the active constraint is removed. There is a substitution θ under which constraints from the constraint store match the head of the rule and satisfy its guard (written $\models \theta \wedge G$).

    The auxiliary function DropIds removes the identifiers from identified constraints. The function NF normalizes terms in an expression. The function Inst replaces fresh unbound variables in a term by fresh constants to ensure a ground term.
  - The (Prop)agate transition is like the (Simp)lify transition except that the active constraint c is matches a kept constraint. It is therefore kept and remains active. The propagation history is therefore updated (and remains unchanged in all other transitions). The function Ids returns the identifiers of identified constraints.

– The (Next) transition, the active constraint cannot be matched against its occurrence in the rule head. One proceeds to the next occurrence in the rules of the program. This makes sure that rules are tried in the order given in the program.
– The (Drop) transition, if there is no more occurrence to try, removes the active constraint the goal store, but it stays (suspended) in the constraint store. The function OccIDs returns all occurrence numbers from the rules in a program.

Figure not part of this draft

Figure 7: The Parallel Part of the Refined CHRe Semantics for 0-Neighbor Restricted Rule

Localised Parallel Transitions. Figure 7 shows the parallel transitions. They are particularly simple. Transition (Single) says that any sequential transition is a parallel transition. Transition (Concurrent) allows to combine two independent transitions on *non-overlapping* parts of the state into one parallel transition. This means that independent computation steps on localized states can be executed in parallel.

Monotonicity, Soundness and Serializability. In the refined CHRe semantics, monotonicty holds with respect to locations, this means computations can be repeated in any larger context of more locations. Serializability holds in that every parallel CHRe computation can be simulated using sequential CHRe transitions. We also prove soundness the refined CHRe semantics with respect to the abstract CHRe semantics.

9.5 Encoding 1- and *n*-Neighbor Rules in Local Rules

We give an encoding of the more general 1-neighbor restricted rules into local, i.e. 0-neighbor restricted rules. We can do the same for *n*-neighbor restricted rules. In this way, a programmer can use *n*-neighbor rules while the translation generates the necessary communication and synchronization between locations. The encodings are a block-free variation of a two-phase commit consensus protocol between locations.

Two-Phase-Commit Consensus Protocol. The protocol consists of two phases:

- *Commit-Request Phase (Voting Phase).* The coordinator process informs all the participating processes about the transaction and to vote either commit or abort. The processes vote.
- *Commit Phase.* If all processes voted commit, the coordinator performs its part of the transaction, otherwise aborts it. The coordinator notifies all processes. The processes then act or abort locally.

The standard protocol can block if a process waits for a reply. Not so in the variation we use.

Encoding 1-Neighbor Restricted Programs. According to the following scheme, we translate each 1-neighbor restricted rule of the form

```
r : [X]Px, [X]Px', [Y]Py \ [X]Sx, [Y]Sy <=> Gx,Gy | Body.
```

In the head, `Px` are the persistent constraints and `Px'` are the non-persistent constraints. Constraints are *persistent* if they are not removed by any rule in the program. In the guard, `Gx` contains only variables from location `X`. In the rule scheme below, `Ays` contains all variables from the rule head, and `Xs` only the variables from location `x`.

```
 % Commit-Request Phase
 % match and send request to neighbor location
request : [X]Px,[X]Sx ==> Gx | [Y]r_req(Xs).
 % match and send commit to primary location
vote : [Y]Py,[Y]Sy \ [Y]r_req(Xs) <=> Gy | [X]r_vcom(XYs). % if Sx non-e.
vote : [Y]Py,[Y]Sy,  [Y]r_req(Xs) ==> Gy | [X]r_vcom(XYs). % if Sx empty

 % Commit Phase
```

```
 % remove non-persistent constraints at primary location and send commit
commit :  [X]Px \ [X]Px',[X]Sx, [X]r_vcom(XYs) <=> [Y]r_commit(XYs).
 % remove at neighbor location, add non-persistent and body constraints
act   : [Y]Py \ [Y]Sy, [Y]r_commit(XYs) <=> [X]Px', Body.
 % otherwise abort, re-introduce removed constraints at primary location
abort : [Y]r_commit(XYs) <=> [X]Px', [X]Sx.
```

The rule scheme uses different `vote` rules depending on the emptiness of `Sx`. If `Sx` is empty, it should be possible to remove several instances of `Sy` with the same request. Note that the rule scheme requires a refined semantics where rules are tried in the given order, because we have to make sure that rule `act` is tried before the abort rule `abort`.

The rule scheme implements an *asynchronous and optimistic consensus protocol* between two locations of the ensemble. It is asynchronous because neither primary nor neighbor location ever block or busy-wait for responses. Rather they communicate asynchronously via the protocol constraints, while potentially interleaving with other computations. The temporary removal of non-persistent constraints in the rule scheme ensures that the protocol cannot be interfered with. It is optimistic because non-protocol constraints are only removed after both locations have independently observed their part of the rule head instance. It is possible that some protocol constraints are left if the transaction did not commit, but these can be garbage-collected.

We can generalize the above encoding to n-neighbor restricted rules.

Quiescence, Soundness and Serializability. We say that a CHRe program is *locally quiescent (terminating)* if given a reachable state, we cannot have any infinite computation sequences that do not include the (Flush) transition. Hence local quiescence guarantees that each location will eventually process the constraints delivered to its buffer.

Serializability and soundness of the encoding for quiescent programs are proven: computations between commit-free states of 0-neighbor restricted encodings have a mapping to computations of the original 1-neighbor restricted program.

Comingle. This new programming language can be described as an extension of CHRe for distributed logic programming over mobile devices [LCF15]. There is a prototype on the Android operating system, see `https://github.com/sllam/comingle`. One application was built both using CoMingle and by writing traditional code: the former was about one tenth of the size of the latter without a noticeable performance penalty.

10 Models of Concurrency in CHR

Theoretical and practical models of concurrency have been encoded in CHR. Such an effective and declarative embedding holds many promises: It makes theoretical models executable. It can serve as executable specification of the practical models. One can toy with alternative design choices. The implementations can be formally verified and analyzed using standard and novel CHR analysis techniques. Last but not least it allows to compare different models on a common basis.

We will shortly introduce some common models of concurrency by their implementation in CHR: Software Transactional Memory, Colored Petri Nets, Actors and Join-Calculus. Typically soundness and completeness results will prove the correctness of these embeddings.

10.1 Software Transactional Memory STM

While we already seen the use of STM to implement parallel CHR in Haskell in section 6. Now we do it the other way round. The paper [SC08] gives a rule-based specification of Haskell's Software Transactional Memory in parallel CHR which naturally supports the concurrent execution of transactions.

We classify CHR constraints into operation constraints and data constraints. We assume CHR rules where the head contains exactly one operation constraint and the body contains at most one operation constraint.

Shared Memory Operations. We first model shared memory and its associated read and write operations in CHR.

```
read  : cell(L,V1) \ read(L,V2) <=> V1=V2.
write : cell(L,V2), write(L,V2) <=> cell(L,V2).
```

L is a location identifier and V1 and V2 are values. cell is a data constraint, read and write are operation constraints. The write rule performs an destructive assignment to update the value of the cell. With indexing and in-place constraint updates, the compiled rule can run in constant time.

STM run-time manager in CHR. The effects of an STM transaction are reads and writes to shared memory. The STM run-time must guarantee that all reads and writes within a transaction happen logically at once. In case transactions are optimistically executed in parallel the STM run-time must take care of any potential read/write conflicts. The STM run-time must ensure that in case of conflicts at least one transaction can successfully commit its updates whereas the other transaction is retried.

To accomplish this behavior, we use for each transaction a read log and a write log. Before we can commit the write log and actually update the memory cell, we first must validate that for each cell whose value is stored in the read log, the actual value is still the same.

In Figure 8, we specify the STM manager via CHR rules. It has been slightly simplified in this survey. Besides locations and values, we introduce an identifier for transactions T. The operation constraints are read and write and the protocol constraints are validate, commit and rollback, retry. The data constraint CommitRight acts as a token a committing transaction has to acquire in order to avoid concurrent writes. The constraint validate is issued at an end of the transaction if the CommitRight is available. Rules for rollback and retry of transactions are not shown here for space reasons.

```
 % Execution phase ---
 % Read from write or read log, create read log otherwise
r1 : WLog(t,l,v1) \ Read(t,l,v2) <=> v1=v2.
r2 : RLog(t,l,v1) \ Read(t,l,v2) <=> v1=v2.
r3 : Cell(l,v1) \ Read(t,l,v2) <=> v1=v2, RLog(t,l,v1).

 % Write to write log, create write log otherwise
w1 : WLog(t,l,v1), Write(t,l,v2) <=> WLog(t,l,v2).
w2 : Write(t,l,v) <=> WLog(t,l,v).

 % Validation phase ---
 % Check and remove read log, rollback on read log conflict
v1 : Cell(l,v1), Validate(t) \ RLog(t,l,v2) <=> v1=v2 | True.
v2 : Cell(l,v1) \ Validate(t), RLog(t,l,v2) <=> v1=\=v2 | Rollback(t).

 % Start commit phase by acquiring CommitRight, otherwise rollback
s1 : CommitRight, Validate(t) <=> Commit(t).
s2 : Validate(t) <=> Rollback(t).

 % Commit phase ---
 % write update cells, then return CommitRight
c1 : Commit(t) \ Cell(l,v1), WLog(t,l,v2) <=> Cell(l,v2).
c2 : Commit(t) <=> CommitRight.
```

Figure 8: STM Run-Time Manager in CHR

Soundness and Correctness. Our implementation guarantees atomicity, isolation and optimistic concurrency. It is therefore sound. It is correct: if a transaction commits successfully, the store reflects correctly all the reads/writes performed by that transaction.

10.2 Colored Petri Nets CPN

Petri nets are diagrammatic formalism to describe and reason about concurrent processes. They consist of labelled *places* (◯) in which *tokens* (•) reside. Tokens can move along *arcs* passing through *transitions* (▭) from one place to another. A transition may have several incoming arcs and several outgoing arcs. A transition can only fire if all incoming arcs present a token. On firing, all incoming tokens will be removed and a token will be presented on each outgoing arc. *Colored Petri Nets (CPN)* significantly generalize Petri nets. Tokens are colored and places are typed by the colors they allow. Transitions can have conditions on tokens and equations that compute new tokens from old ones.

The paper [Bet07] shows that (Colored) Petri nets can easily be embedded into CHR. When CPNs are translated to CHR, color tokens are encoded as numbers. Place labels are mapped to CHR constraint symbols, tokens at a place to instances of CHR constraints, transitions and their arcs to simplification rules. Incoming arc places form the rule head, outgoing arc places form the rule body, and the transition conditions as well as equations form the rule guard.

Example 10.1. For simplicity, we consider the dining philosophers problem with just three philosophers as CPN in Figure 9. Each philosopher (and fork) corresponds to a colored token, given as a number from 0 to 2. Two philosophers x and y are neighboring if $y = (x+1)$ *mod* 3. Places are think, eat and fork, transitions are eat-think and and think-eat.

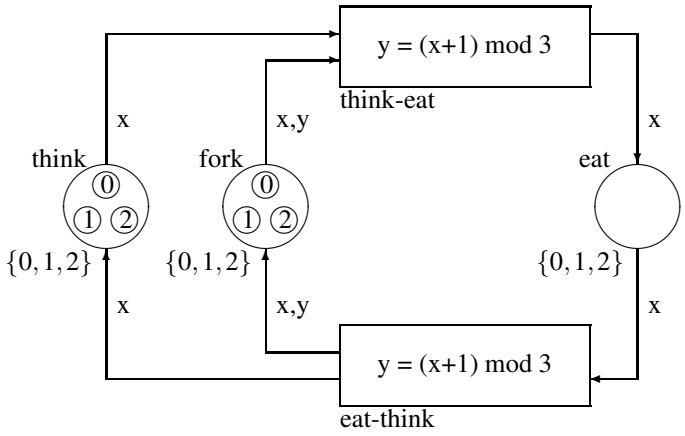

Figure 9: The Three Dining Philosophers Problem as Colored Petri Net

The CPN of Figure 9 translates into the following two CHR rules

```
think_eat : think(X), fork(X), fork(Y) <=> Y =:= (X+1) mod n | eat(X).
eat_think : eat(X) <=> Y =:= (X+1) mod n | think(X), fork(X), fork(Y).
```

Soundness and Completeness. For both classical and Colored Petri nets, these correctness theorems are proven for the translation into CHR.

10.3 Actor Model

In the actor model, one coordinates concurrent computations by message passing. Actors communicate by sending and receiving messages. Sending is a non-blocking asynchronous operation. Each sent message is placed in the actors mailbox (a message queue). Messages are processed via receive clauses which perform pattern matching ad

guard checks. Receive clauses are tried in sequential order. The receive operation is blocking. If none of the receive clauses applies the actor suspend until a matching message is delivered. Receive clauses are typically restricted to a single-headed message pattern. That is, each receive pattern matches at most one message.

In [SLVW08], we extend the actor model with receive clauses allowing for multi-headed message patterns. Their semantics is inspired by their translation into CHR. We have implemented a prototype in Haskell `https://code.google.com/archive/p/haskellactor/`.

Example 10.2. In the Santa Clause problem, Santa sleeps until woken by either all of his nine reindeer or by three of his ten elves. If woken by the reindeer, he harnesses each of them to his sleigh, delivers toys and finally unharnesses them. If woken by three elves, he shows them into his study, consults with them on toys and finally shows them out. Here is a solution using the proposed multi-head extension:

```
santa sanActor =
 receive sanActor of
   Deer x1, Deer x2, ..., Deer x8, Deer x9 -> harness, deliver, unharness.
   Elf x1, Elf x2, Elf x3 -> enter_study, consult, leave_study.
```

This straightforward solution avoids the clumsiness of explicitly counting deers and elfs in the mailbox. There is an obvious direct embedding of the matching receive clauses into CHR simplification rules.

Semantics of Actors with Multi-Headed Message Patterns. We study two possible semantics for this extension, inspired by the standard refined semantics of CHR:

- The *first-match semantics* provides a conservative extension of the semantics of single-headed receive clauses. This semantics guarantees monotonicity: any successful match remains valid if further messages arrive in the actors mailbox.
- The *rule-order-match semantics* guarantees that rule patterns are executed in textual order. In this semantics, newly arrived messages can invalidate earlier match choices.

It will depend on the application which semantics is the better choice.

10.4 Join-Calculus and Join-Patterns

In Join-Calculus, concurrency is expressed via multi-headed declarative reaction rules that rewrite processes or events. The (left-hand side of a) rule is called *join-pattern*. In the thesis [Lam11b] the author extends join-patterns with guards and describes a prototype implementation in parallel CHR compiled to Haskell, see `http://code.haskell.org/parallel-join`.

Join-Calculus with Guarded Join-Patterns. A concurrent *process (or event)*, say `P`, has the form of a predicate. We introduce *guards* into these rules:

```
Guarded Reaction Rule     P1,...Pn if Guard => P1',...Pm'
```

The Join-Calculus semantics is defined by a chemical abstract machine (CHAM). This model specifies transformations using a chemical reaction metaphor. The CHAM can be embedded in CHR, see chapter 6 in [Frü09].

Example 10.3. A print job is to be executed on any available printer where it fits. So print jobs have a size, and printers have a certain amount of free memory. This behavior is captured by the following guarded reaction rule:

```
ReadyPrinter(p,m), Job(j,s) if m>s => SendJob(p,j)
```

There is an obvious direct translation into CHR simplification rules.

Implementation and Experimental Results. CHR goal-based lazy matching is a suitable model for computing the triggering of join-patterns with guards: each process (CHR goal) essentially computes only its own rule head matches asynchronously and then proceeds immediately. We conducted experiments of our parallel Join-Calculus implementation with examples for common parallel programming problems. They show consistent speed-up as we increase the number of processors.

11 Conclusions

We have given an exhaustive survey of abstract and more refined semantics for parallel CHR as well as distributed CHR. Most of them have been proven correct. These semantics come with several implementations in both software and hardware. We shortly presented a few applications and classical example programs and promising experimental results. Last but not least we reviewed concurrency models that have been encoded in CHR to get a better understanding of them and sometimes to extend them. Most embeddings have been proven sound and complete. Some implementations and embeddings are available online.

We now present common topics and issues that we have identified as a result of this survey and that lead to research questions for future work.

Syntactic Fragments of CHR. The parallel semantics surveyed are concerned with expressive Turing-complete fragments of CHR. They typically consider the subclass of *Ground CHR without propagation rules*. Hardware implementations rely on *non-size-increasing rules* which are still Turing complete.

Sometimes the notion of constraints is too abstract, and one differentiates between *data and operation constraints*. This clarifies programs like Blockworld and Union-Find, is essential in the semantics of CHR with transactions and in the concurrency model of Software Transactional Memory when encoded in CHR.

Propagation Rules. Surprisingly, while propagation rules seem perfect for parallelization, because they do not remove any constraints, they are currently only supported in distributed CHRd and CHRe. On the other hand it seems only a matter of technicality to extend the refined parallel semantics with propagation rules, either using the *propagation history* of CHRe or the *occurrence check* approach of CHRd to avoid their trivial non-termination. The latter could also be applied to CHRmp.

Propagation rules seem to favor a set-based semantics if occurrence checks are used as in CHRd. The ideas from the *persistent constraint semantics* of CHR [BRF10b] could make propagation rules more steadfast in these semantics. In this semantics, propagation rules produce constraints that cannot be removed, i.e. persistent ones. This greatly improves the behavior of the occurrence check.

Serializability and Soundness. These properties have been proven for all parallel CHR semantics based on multisets, *monotonicity* for most. Surprisingly, these properties do not hold in full for the set-based semantics of distributed CHRd and massively-parallel CHRmp, though CHRmp is sound for *deletion-acyclic programs*.

Program Analysis. We should re-examine CHR program analysis for parallel CHR semantics to see how they carry over. *Termination* is quiescence in the concurrent context.

Confluence is an essential desirable property of sequential CHR programs. It already plays a role in parallel CHR for sound removal of transactions and seems trivial in exhaustively-parallel CHRmp. Confluence seems strongly related to soundness and serializability properties of concurrent CHR semantics. Semi-automatic *completion* can generate rules to to of non-confluent programs. This method has been used in parallelizing the Union-Find algorithm and can be used in translating away CHR transactions.

We think that the notion of *deletion-acyclicity* has a broader application in rule-based systems. It seems related to confluence and can be expressed as a termination problem.

Implementations. A full-fledged widely available stable implementation of parallel CHR is still missing. It could serve as a basis to foster further research and applications, as does the K.U. Leuven platform for sequential CHR. With Comingle, the situation seems better in the case of distributed CHR.

Models of Concurrency. Embedding models of concurrency in CHR is promising for understanding, analyzing and extending models, but still in its infancy. It is appealing because of the *lingua franca* argument for CHR: different embeddings can be compared on its common basis and fertilize each other. Conversely, the striking similarity of some models when encoded in CHR leads one to speculate about a generic concurrency model that is a suitable fragment of CHR which could then be mapped to many existing models, yielding a truly unified approach.

Last but not least, it should be investigated how the research surveyed here carries over to related languages like constraint logic programming ones and the other rule-based approaches that have been embedded in CHR. In general, the CHR research surveyed here should be related to more mainstream research in concurrency, parallelism and distribution.

References

[AFM99] Slim Abdennadher, Thom Frühwirth, and Holger Meuss. Confluence and Semantics of Constraint Simplification Rules. *Constraints*, 4(2):133–165, 1999.

[Bet07] Hariolf Betz. Relating coloured Petri nets to Constraint Handling Rules. In Djelloul et al. [DDS07], pages 33–47.

[Bet14] Hariolf Betz. *A Unified Analytical Foundation for Constraint Handling Rules*. BoD–Books on Demand, 2014.

[BRF10a] Hariolf Betz, Frank Raiser, and Thom Frühwirth. A complete and terminating execution model for constraint handling rules. *Theory and Practice of Logic Programming*, 10:597–610, 7 2010.

[BRF10b] Hariolf Betz, Frank Raiser, and Thom Frühwirth. Persistent constraints in Constraint Handling Rules. In A. Wolf and U. Geske, editors, *WLP '09: Proc. 23rd Workshop on (Constraint) Logic Programming*. Universität Potsdam, 2010.

[DDS07] K. Djelloul, G. J. Duck, and M. Sulzmann, editors. *CHR '07: Proc. 4th Workshop on Constraint Handling Rules*, September 2007.

[DSGH04] Gregory J. Duck, Peter J. Stuckey, María García de la Banda, and Christian Holzbaur. The refined operational semantics of Constraint Handling Rules. In B. Demoen and V. Lifschitz, editors, *ICLP '04*, volume 3132 of *LNCS*, pages 90–104. Springer, September 2004.

[FR11] Thom Frühwirth and Frank Raiser, editors. *Constraint Handling Rules: Compilation, Execution, and Analysis*. Books on Demand, March 2011.

[Frü05a] Thom Frühwirth. Parallelizing union-find in Constraint Handling Rules using confluence. In M. Gabbrielli and G. Gupta, editors, *ICLP '05*, volume 3668 of *LNCS*, pages 113–127. Springer, October 2005.

[Frü05b] Thom Frühwirth. Specialization of concurrent guarded multi-set transformation rules. In S. Etalle, editor, *LOPSTR '04*, volume 3573 of *LNCS*, pages 133–148. Springer, 2005.

[Frü06] Thom Frühwirth. Deriving linear-time algorithms from union-find in CHR. In T. Schrijvers and Th. Frühwirth, editors, *CHR '06*, pages 49–60. K.U.Leuven, Dept. Comp. Sc., Technical report CW 452, July 2006.

[Frü09] Thom Frühwirth. *Constraint Handling Rules (Monography)*. Cambridge University Press, August 2009.

[Frü15] Thom Frühwirth. Constraint handling rules – what else? In *Rule Technologies: Foundations, Tools, and Applications*, pages 13–34. Springer International Publishing, 2015.

[Frü16] Thom Frühwirth. *The CHR Web Site* – `www.constraint-handling-rules.org`. Ulm University, 2016.

[GdlBP08] M. García de la Banda and E. Pontelli, editors. *ICLP '08: Proc. 24rd Intl. Conf. Logic Programming*, volume 5366 of *LNCS*. Springer, December 2008.

[GM09] Maurizio Gabbrielli and Maria Chiara Meo. A compositional semantics for CHR. *ACM TOCL*, 10(2):1–36, February 2009.

[GMTW13] Maurizio Gabbrielli, Maria Chiara Meo, Paolo Tacchella, and Herbert Wiklicky. Unfolding for CHR programs. *Theory and Practice of Logic Programming*, pages 1–48, 2013.

[Hae11] Rémy Haemmerlé. (Co-)Inductive semantics for Constraint Handling Rules. volume 11(4–5) of *TPLP*, pages 593–609. Cambridge University Press, July 2011.

[HGSD05] Christian Holzbaur, María García de la Banda, Peter J. Stuckey, and Gregory J. Duck. Optimizing compilation of Constraint Handling Rules in HAL. volume 5(4–5) of *Theory and Practice of Logic Programming*, pages 503–531. Cambridge University Press, July 2005.

[Lam11a] Edmund S.L. Lam. Concurrent chr, chapter 5. In Thom Frühwirth and Frank Raiser, editors, *Constraint Handling Rules: Compilation, Execution, and Analysis*. Books on Demand, March 2011.

[Lam11b] Edmund Soon Lee Lam. *Parallel execution of Constraint Handling Rules – Theory, Implementation and Application*. PhD thesis, School of Computing, Department of Computing Science, National University of Singapore, 2011.

[LC13] Edmund S.L. Lam and Iliano Cervesato. Decentralized execution of constraint handling rules for ensembles. In *Proceedings of the 15th Symposium on Principles and Practice of Declarative Programming*, pages 205–216. ACM, 2013.

[LCF15] Edmund Soon Lee Lam, Iliano Cervesato, and Nabeeha Fatima. Comingle: Distributed logic programming for decentralized mobile ensembles. In *Coordination Models and Languages - 17th IFIP WG 6.1 International Conference, COORDINATION 2015, Grenoble, France, 2015*, pages 51–66, 2015.

[LS07] Edmund S.L. Lam and Martin Sulzmann. A concurrent Constraint Handling Rules semantics and its implementation with software transactional memory. In *DAMP '07: Proc. ACM SIGPLAN Workshop on Declarative Aspects of Multicore Programming*. ACM Press, January 2007.

[LS09] Edmund S.L. Lam and Martin Sulzmann. Concurrent goal-based execution of Constraint Handling Rules. *TPLP*, 11:841–879, 2009.

[Mei07] Marc Meister. Concurrency of the preflow-push algorithm in Constraint Handling Rules. In *CSCLP'07: Proc. 12th Intl. Workshop on Constraint Solving and Constraint Logic Programming*, pages 160–169, 2007.

[MF07] Marc Meister and Thom Frühwirth. Reconstructing almost-linear tree equation solving algorithms in CHR. In *Proceedings of CSCLP 2007: Annual ERCIM Workshop on Constraint Solving and Constraint Logic Programming*, page 123, 2007.

[RBF09] Frank Raiser, Hariolf Betz, and Thom Frühwirth. Equivalence of CHR states revisited. In F. Raiser and J. Sneyers, editors, *CHR '09*, pages 33–48. K.U.Leuven, Dept. Comp. Sc., Technical report CW 555, July 2009.

[RF10] Frank Raiser and Thom Frühwirth. Exhaustive parallel rewriting with multiple removals. In Slim Abdennadher, editor, *WLP '10*, September 2010.

[SC08] Martin Sulzmann and Duc Hiep Chu. A rule-based specification of Software Transactional Memory. In M. Hanus, editor, *LOPSTR '08, Pre-proceedings*, 2008.

[SL07] Martin Sulzmann and Edmund S.L. Lam. Compiling Constraint Handling Rules with lazy and concurrent search techniques. In Djelloul et al. [DDS07], pages 139–149.

[SL08] Martin Sulzmann and Edmund S.L. Lam. Parallel execution of multi-set constraint rewrite rules. In S. Antoy, editor, *PPDP '08: Proc. 10th Intl. Conf. Princ. Pract. Declarative Programming*, pages 20–31. ACM Press, July 2008.

[SLVW08] Martin Sulzmann, Edmund S.L. Lam, and Peter Van Weert. Actors with multi-headed message receive patterns. In D. Lea and G. Zavattaro, editors, *COORDINATION '08: Proc. 10th Intl. Conf. Coordination Models and Languages*, number 5052 in LNCS, pages 315–330. Springer, May 2008.

[Sne08] Jon Sneyers. Turing-complete subclasses of CHR. In García de la Banda and Pontelli [GdlBP08], pages 759–763.

[SS08] Tom Schrijvers and Martin Sulzmann. Transactions in Constraint Handling Rules. In García de la Banda and Pontelli [GdlBP08], pages 516–530.

[SSR07] Beata Sarna-Starosta and C.R. Ramakrishnan. Compiling Constraint Handling Rules for efficient tabled evaluation. In M. Hanus, editor, *PADL '07: Proc. 9th Intl. Symp. Practical Aspects of Declarative Languages*, volume 4354 of *LNCS*, pages 170–184. Springer, January 2007.

[SSSD07] Beata Sarna-Starosta, R. E. Kurt Stirewalt, and Laura K. Dillon. A model-based design-for-verification approach to checking for deadlock in multi-threaded applications. *Intl. Journal of Softw. Engin. and Knowl. Engin.*, 17(2):207–230, 2007.

[SVWSDK10] Jon Sneyers, Peter Van Weert, Tom Schrijvers, and Leslie De Koninck. As time goes by: Constraint Handling Rules – A survey of CHR research between 1998 and 2007. *TPLP*, 10(1):1–47, 2010.

[TORF12] Andrea Triossi, Salvatore Orlando, Alessandra Raffaetà, and Thom Frühwirth. Compiling chr to parallel hardware. In *Proceedings of the 14th symposium on Principles and practice of declarative programming*, pages 173–184. ACM, 2012.

[VW10] Peter Van Weert. Efficient lazy evaluation of rule-based programs. *IEEE Transactions on Knowledge and Data Engineering*, 22(11):1521–1534, November 2010.

[ZFG12] Amira Zaki, Thom Frühwirth, and Ilvar Geller. Parallel execution of Constraint Handling Rules on a Graphical Processing Unit. In Jon Sneyers and Thom Frühwirth, editors, *CHR '12*, pages 82–90. K.U.Leuven, Dept. Comp. Sc., Technical report CW 624, September 2012.